Measuring the Phase Variation of a DOCSIS 3.1 Full Duplex Channel

A Thesis Submitted
to the College of Graduate and Postdoctoral Studies
in Partial Fulfillment of the Requirements
for the Degree of Master of Science
in the Department of Electrical and Computer Engineering
University of Saskatchewan

by

Sean M. Wilson

Saskatoon, Saskatchewan, Canada

Permission to Use

In presenting this thesis in partial fulfillment of the requirements for a Postgraduate degree from the University of Saskatchewan, it is agreed that the Libraries of this University may make it freely available for inspection. Permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes may be granted by the professors who supervised this thesis work or, in their absence, by the Head of the Department of Electrical and Computer Engineering or the Dean of the College of Graduate Studies and Research at the University of Saskatchewan. Any copying, publication, or use of this thesis, or parts thereof, for financial gain without the written permission of the author is strictly prohibited. Proper recognition shall be given to the author and to the University of Saskatchewan in any scholarly use which may be made of any material in this thesis.

Request for permission to copy or to make any other use of material in this thesis in whole or in part should be addressed to:

Head of the Department of Electrical and Computer Engineering
57 Campus Drive
University of Saskatchewan
Saskatoon, Saskatchewan, S7N 5A9, Canada

OR

Dean
College of Graduate and Postdoctoral Studies
University of Saskatchewan
116 Thorvaldson Building, 110 Science Place
Saskatoon, Saskatchewan, S7N 5C9, Canada

Abstract

Including a Full Duplex option into DOCSIS introduces several problems. One of the more troublesome issues is the presence of a strong self interference signal that leaks from the transmit side to the receive side of a cable node. This self interference is caused by echoes in the channel that translate the forward travelling transmit signals into a reverse travelling signal, as well as, by leakage from the hybrid coupler used to couple the upstream and downstream signals. To suppress this self interference an echo canceller is implemented to remove the unwanted interference from the received signal. Unfortunately with the high rates of data transmission used in modern day CATV networks the echo canceller needs tremendous precision.

A major concern in the implementation of Full Duplex into DOCSIS is if the channels used are even very slightly time varying. The echos in such channels change with time and can be difficult for the echo canceller to track. Changes in the response of the channel cause the echo profile of the network to shift and the echo canceler to re-adapt to the new channel response. The issue with this changing response is that it is possible for the channel to change faster than the echo canceller can adapt, resulting in the interference becoming unacceptably high. Since the channel is a physical network of coaxial cables often exposed to the environment, its propagation properties can be affected by wind swaying pole mounted cables, or by rapid heating from the sun, or sudden shifts in the load of the network. With information on how the physical properties of the cable changes, the engineers designing the echo canceller can know how fast the canceller must adapt to changes and also have a better measure of how reliable its echo cancellation will be.

In this thesis the stability of the echo profile of the channel is measured. It is shown that the property of the channel with the greatest potential to rapidly change and cause noise after echo cancellation is the phase response of the channel. Due to this, the approach of this thesis is to measure the fluctuations in the phase of the channel response of a CATV network constructed in the lab. To measure the fluctuations in the phase response of the

channel, a PLL (Phase Locked Loop) based circuit is designed and built on an FPGA (Field Programmable Gate Array) and connected to a model of a simple CATV network. The PLL circuit used to measure the phase fluctuations of the channel is designed to be able to measure changes occurring faster than 0.1 Hz and with a power higher than $10^{-7}\,V^2$. The circuit is able to capture data from the channel over a period of 90 seconds.

Using this phase variation measurement circuit a series of experiments were performed on a model CATV DOCSIS network. It was found that many physical disturbances to the network had the effect of rapidly shifting the phase response of the network. Heating the cables in the network was found to shift the phase response upwards of 20000 μradians. Flexing the cables in the network was found to have a peak phase variation of 8000 μradians with similar effects found from walking over cables. Overall, it was clear that physical effects on the network had the propensity to rapidly shift the network response. Any echo canceller that is designed in the future will have to consider these effects when reporting the cancellation that it is able to achieve.

Acknowledgments

I would like to thank my supervisor Dr. Eric Salt for all of his help and guidance throughout the project. I would also like to thank Dr. Brian Berscheid for his expertise and assistance with many parts of my project. I would like to thank Vecima Networks for providing the opportunity to research some practical topics as well as their support throughout my research.

I greatly appreciate my parents putting up with me and giving me somewhere to live through out my master's degree. I would also like to thank my grandparents for always giving me support and motivation to succeed.

Finally I would like to thank my colleagues, Lance Pitka and Quang Nguyen, for helping me stay sane during the writing of this thesis.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

ADC	Analog Digital Converter
AM	Amplitude Modulation
AWGN	Additive White Gaussian Noise
BCH-LDPC	Bose-Chaudhuri-Hocquemghemv Low Density Parity Check
BER	Bit Error Rate
BW	Bandwidth
CATV	Community Antenna TeleVision
CM	Cable Modem
CMTS	Cable Modem Termination System
CNR	Carrier Noise Ratio
CP	Cyclic Prefix
DAC	Digital Analog Converter
DFT	Discrete Fourier Transform
DS	Down Stream
DOCSIS	Data Over Cable System Interface Specification
DSL	Digital Subscriber Line
ENOB	Effective Number Of Bits
FDX	Full Duplex
FEC	Forward Error Correction
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
FPGA	Field Programmable Gate Array
HDX	Half Duplex
HFC	Hybrid Fibre Coaxial
I	Inphase
IDFT	Inverse Discrete Fourier Transform

IFFT	Inverse Fast Fourier Transform
ISI	Inter-Symbol Interference
LMS	Least Mean Squared
LPF	Low Pass Filter
MAX	Media Access Control
MER	Modulation Error Ratio
NCO	Numerically Controlled Oscillator
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
PA	Power Amplifier
PAPR	Peak to Average Power Ratio
PDF	Probability Density Function
PHY	Physical Layer
PLC	Phy Link Layer
PLL	Phase Locked Loop
PSD	Power Spectral Density
Q	Quadrature
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
RAM	Random Access Memory
RLS	Recursive Least Squared
RMS	Root Mean Squared
ROM	Read Only Memory
RP	Roll-Off Period
RX	Receiver
SNR	Signal to Noise Ratio
SOI	Signal Of Interest
TX	Transmitter

US	Up Stream
VNA	Vector Network Analyzer

List of Symbols

A_{ref}	Amplitude of sinusoid being input to the PLL
$\boldsymbol{a}(t)$	Time varying amplitude of a sinusoid
α_r	Attenuation constant associated with an echo at point r
$\delta[n]$	Discrete impulse sequence
$\delta(t)$	Continuous impulse signal
$\Delta\boldsymbol{t}[n]$	Sequence of time jitter samples in seconds
$\Delta\boldsymbol{\phi}[n]$	Sequence of time jitter samples in radians
$e[n]$	Difference signal of an adaptive filter
f	Digital frequency in cycles/sample
F	Analog frequency in hertz
F_s	Sampling frequency in hertz
F_{PLL}	Sampling rate of the PLL
$\mathcal{F}\{x[n]\}$	DOCSIS FFT of $x[n]$
$\mathcal{F}^{-1}\{X[k]\}$	DOCSIS IFFT of $X[k]$
$G(z)$	Open loop PLL transfer function
$h[n]$	Discrete system impulse response
$h(t)$	Time domain channel response
$h_{ij}(t)$	Response between points i and j
$H(\Omega)$	Analog frequency domain response
$H(z)$	Closed loop PLL transfer function
$H_{lf}(z)$	Transfer function of PLL loop filter
$H_{notch}(z)$	Tranfer function of a notch filter
K_x	Gain of scaling factor "x" in a PLL
K_d	Phase detector gain
n	Sample number
$n(t)$	Time domain thermal noise

N_{FFT}	FFT/DFT size
N_{CP}	Cyclic Prefix size
N_{RP}	Roll off Period size
N_r	Number of bits used to represent the frequency in an NCO
N_a	Number of bits used to represent the phase in an NCO
N_d	Number of bits used to represent output sinusoid in an NCO
$\phi[n]$	Discrete time varying phase in cycles
$\phi^r[n]$	Discrete time varying phase in radians
$\phi_i[n]$	PLL input phase in cycles
$\phi_o[n]$	PLL output phase in cycles
$\phi_e[n]$	PLL phase difference in cycles
$\phi_{det}[n]$	Output of phase detector in a PLL in cycles
$\phi(t)$	Continuous time varying phase in radians
P_x	Digital power of signal x in V^2
$q_x[n]$	Quantization noise sequence introduced by x
ρ	Magnitude of a pole on the unit circle
$\widetilde{r}(t)$	Complex passband signal received at the receive side of the hybrid including distortions introduced by network
$R_1(z)$	PLL scaling factor = $K_d \cdot H_{lf}(z) \cdot K_1$
σ_x^2	Variance of x
s	Step size of a discrete signal
$S_{xx}(2\pi F)$	Analog power spectral density of x
$S_{xx}(e^{j2\pi f})$	Digital power spectral density of x
t_r	Time a echo at point r takes to propagate to the received side of the hybrid
T_c	Coherence time of the channel in seconds
T_d	Dwell time of an echo canceller in seconds
μ_x	Mean of x

V	Unit for a digital signal relating electric potential of the input of ADCs and output of DACs to the digital signal level
$V^2 samples$	Digital energy, samples are analogous to seconds in the digital domain so the energy is the sum of the squares of digital samples
V^2	Digital power, average power in a sequence of digital samples is the digital energy divided by the number of samples averaged by
$\frac{V^2}{cycle/sample}$	Digital power spectral density. V^2 per digital cycle related to sampling frequency.
$w[n]$	Taps of a digital filter
ω	Digital frequency in radians
Ω	Analog frequency in radians
$\overline{x^2[n]}$	Mean squared of x[n]
$\widetilde{x}[n]$	Real passband version of complex baseband x[n]
$\widetilde{x_1}(t)$	Complex passband signal transmitted by the node
$\widetilde{x_2}(t)$	Complexx passband signal transmitted by the cable modems
$y[n]$	Output of a discrete time system
z	Result of Z transform $= e^{jw}$

1. Introduction

1.1 History of Cable Systems

The history of CATV (Common Antenna Television) or cable networks probably began in the late 1940's when a radio operator named L.E. "Ed" Parsons responded to a challenge from his wife who wanted pictures with her radio [1]. Ed went to the tallest nearby building and constructed an antenna able to pick up a weak television signal which he was able to connect to his television. Before long, much of his town wanted to be connected to this antenna to receive a television signal and the CATV boom had started. When Ed was able to pick up the television signal, he was responding to a need felt through much of rural America. These Americans were too distant from most commercial stations to pick up a television signal with a regular antenna. In other parts of America, CATV was simultaneously developing and consumers soon clamored to be connected to a cable network [1]. This quickly resulted in CATV becoming large business and the technology slowly began to develop into what we see today.

In the next few decades several technological improvements began to make their way into the traditional CATV networks. The improvements mainly consisted of increasing the number of channels that the cable network was able to support. An important technological innovation was the introduction of a upstream path in CATV networks. The upstream path would allow cable subscribers to communicate back to the head end of the CATV network and thus two-way communications was achievable. The upstream communication technolody was originally developed for military applications, but was quickly integrated

into commercial CATV which led to the birth of early cable modems [2].

The rapid growth of CATV networks continued up until the late 1980's with the advent of the internet. It was around this time that cable providers switched from a traditional analog distribution to a digital system. A key part of this switch was to divide the existing network into HFC (Hybrid Fibre Coaxial) networks which allowed for faster data transmission. By making these changes cable networks were able to compete with phone companies which had begun to offer services such as DSL (Digital Subscriber Line), which offered much faster data transmission than cable was initially able to provide [3].

The CATV network became digital and consumers modern demands were now able to be met. CATV subscribers wanted modern services such as pay TV, point-to-point services such as teleconferencing and high-speed data exchange, and eventually the internet. The main drive for this advancement was pay television such as Home Box-Office (HBO). The cable provider's cable modems that they had located in the subscribers homes slowly transformed into devices that were able to provide all of these services.

Early on the electronics for cable networks was proprietary equipment from a single manufacturer [4]. While there was more then one equipment manufacturer, a cable network had to commit to one, which meant that manufactures had a monopoly on any expansion to the network. In order to create competition among manufacturers, service providers formed a not-for-profit organization in 1988. The purpose of this organization, which was called CableLabs, is to create standards that would allow interoperability of equipment supplied by different manufactures [5]. CableLabs introduced its first interoperable standard, referred to as Data Over Cable Service Interface Specification (DOCSIS), in 1997 [6].

1.2 Modern Cable Systems

The first version of the North American data over cable standard, known as DOCSIS 1.0, was released by Cable Labs in 1997. In this initial version of DOCSIS, the DOCSIS network was described as many Cable Modems (CMs) located in subscriber households connected

to a Cable Modem Termination System (CMTS) via a HFC network [6]. In the upstream direction (from the households to the head end,) the CMs were described as transmitters and the CMTS as a receiver. In the downstream direction (from the head end to the household), the CMs were treated as receivers and the CMTS as the transmitter. After the release of DOCSIS 1.0, the standard continued to evolve to meet the demands of modern consumers. The demand for more bandwidth, particularly in the upstream direction, was a source of major upgrades to DOCSIS.

Following the releasal of DOCSIS 1.0 in 1997 DOCSIS 1.1, which added VOIP (Voice Over IP) service, was released. DOCSIS 1.0 and DOCSIS 1.1 are collectively referred to as DOCSIS 1.x. Following DOCSIS 1.x was DOCSIS 2.0, which released December 31, 2001 and offer a 3X increase in upstream speeds. DOCSIS 3.0 was released August 4, 2006. This release offered much higher speeds in both directions. Finally DOCSIS 3.1 was released on October 29, 2013 which again offered greatly increased speeds and higher efficiency. The next iteration, known as Full Duplex (FDX) DOCSIS 3.1, was released on October 26, 2017. It promises symmetric gigabit service greatly increasing the upstream capacity of DOCSIS 3.1 [7].

Table 1.1 [7] summarizes the capability of the sequence of DOCSIS standards. It shows the steady increase of data throughput and services available through the standard. The improvements in DOCSIS were mainly due to higher Quadrature Amplitude Modulation (QAM) orders on the upstream and downstream paths, as well as increased support for channel bonding and Internet Protocol Version 6 (IPv6). The increase in the modulation order of the QAM allowed for more throughput over the same bandwidth. With DOCSIS 3.1 the modulation scheme was augmented to include Orthogonal Frequency Division Multiple Access (OFDMA) to further increase data throughput and spectral efficiency. The demand for more bandwidth seems to be ever increasing, which will continue to pressure cable providers to innovate and adapt.

Table 1.1 DOCSIS Standard Progression

	Highlights	**Downstream Capacity**	**Upstream Capacity**	**Release Date**
DOCSIS 1.0	Initial Cable Broadband Technology	40 Mbps	10 Mbps	1997
DOCSIS 1.1	Added VOIP Service	40 Mbps	10 Mbps	2001
DOCSIS 2.0	Higher Upstream Speed	40 Mbps	30 Mbps	2002
DOCSIS 3.0	Greatly Enhanced Capacity	1 Gbps	100 Mbps	2006
DOCSIS 3.1	Capacity and Efficiency Progression	10 Gbps	1-2 Gbps	2013
Full Duplex DOCSIS 3.1	Symmetrical Streaming and Increased Upload Speeds	10 Gbps	10 Gbps	2017

1.3 Future Cable Systems

The demand for high speed internet service has continued to grow and will continue to grow into the foreseeable future. It is currently at a point where it has begun to be economical for telephone companies to plow fibre directly to the homes of subscribers to meet the demand. Cable service providers believe that they can compete with this fibre to the home service without having to completely replace their expensive coaxial cable plants with optical fibres. As described in the previous section the cable service providers have been incrementally increasing the throughput of their cable plants through a series of technological changes. Two proposed changes of particular interest to the cable providers are decentralizing the CMTS by remoting the physical layer, referred to as Remote Phy, and introducing new technology that supports Full Duplex, which allows for simultaneous upstream and downstream communication in the same spectrum.

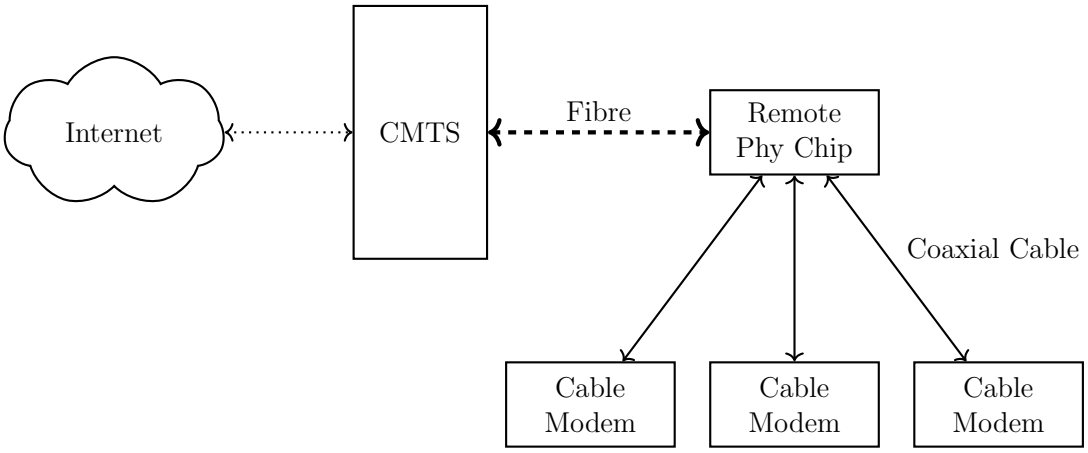

Figure 1.1 Remote Phy CATV System

1.3.1 Remote Phy

For sending information up and down fibre there are at least two fundamental methods in the context of HFC plants: Linear HFC and Digital HFC. Up until recently Linear HFC was most widely used in the cable market due to its relative inexpensiveness and ease of use. However, Linear HFC has some clear disadvantages. Linear HFC is limited in the distance it can send data, it requires significant upkeep, and it introduces noise that can not be removed. Increasingly the industry is moving towards Digital HFC as it can travel long distances, support more wavelengths, has lower fixed and operating costs, scales well, and the noise inherent to the signal can be removed [8]. Digital HFC has the disadvantage of being more complicated than Linear HFC. Changing DOCSIS to Digital HFC can be accomplished by using DOCISIS Remote Phy technology. DOCSIS currently uses a centralized CMTS (Cable Modem Termination System) located at the cable providers head end. This CMTS has a phy chip that modulates data into a signal appropriate for transmission on a coaxial cable network. The basic premise of Remote Phy is to take this phy chip and locate it closer to the subscriber. The CMTS would be connected to the phy chip through fibre lines using Digital HFC.

1.3.2 Full Duplex

A hot topic currently for cable providers is the switch to full duplex, which promises great increases in the upstream communication speeds that can be offered. Historically bandwidth was used for communication in one direction only, and as a result the spectrum is allocated in such a way that the downstream has a much larger slice than the upstream. Today, there is more demand for upstream traffic and more upstream bandwidth must be allocated in order to support the increased bit rates needed for such things as cloud services, uploading pictures, and online games. However, the amount of spectrum available is finite and an increase in upstream capacity cannot come at the cost of reduced downstream. In the old frequency division duplex system (or half duplex), the spectrum was split between the upstream and the downstream with the upstream occupying the lower frequency band and the downstream occupying the upper as shown in Figure 1.2. The next logical step in the evolution of cable systems is to use a full duplex system in which upstream and downstream spectra occupy the same portion of bandwidth at the same time. Switching to full duplex will vastly increases the upstream bandwidth without sacrificing downstream, as shown in Figure 1.3. Moving to a full duplex system would increase the data throughput in the upstream direction and allow the coaxial cable networks to remain competitive with fibre to the home networks for years to come.

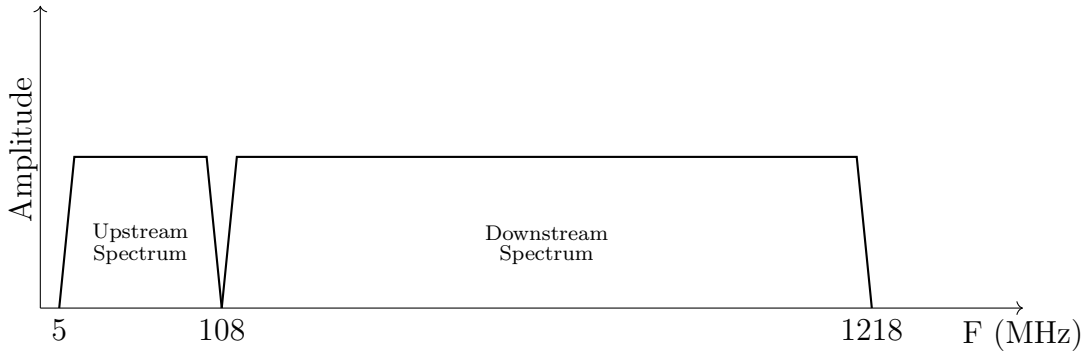

Figure 1.2 Half Duplex Spectrum

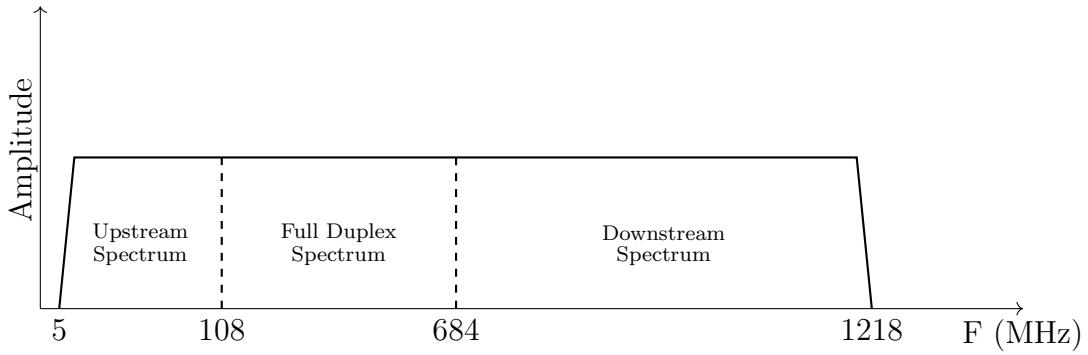

Figure 1.3 Full Duplex Spectrum

1.4 Obstacles in Full Duplex

Changing a network from one that supports frequency division duplex to one that supports a mix of full duplex and frequency division duplex requires a number of changes to the hardware in the system. Perhaps the most significant of which is modifying the diplexer, which is essential to frequency division duplex, to allow two way communication. Modifying the diplexer means replacing it with a specialized directional coupler as shown in Figure 1.4 as a "Hybrid".

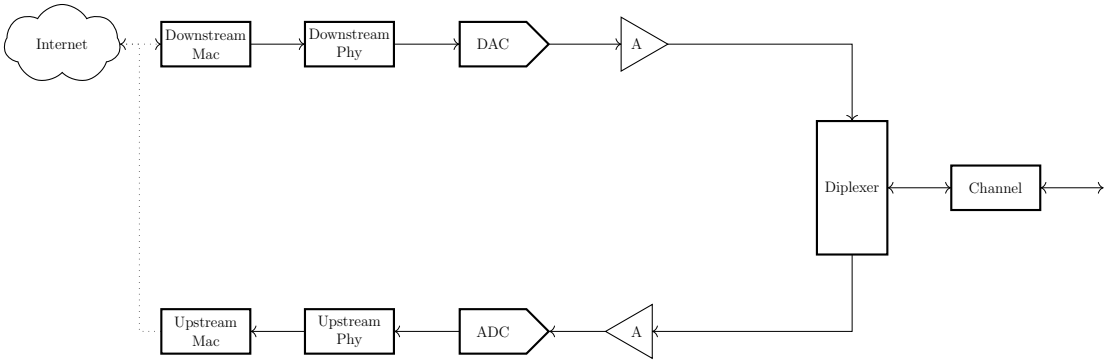

Figure 1.4 Block Diagram of a Full Duplex Node Without an Echo Canceller

A diplexer separates the upstream spectrum from the downstream spectrum using strategically placed low pass and high pass filters. However, the full duplex portion of the system supports two way communication over part of the bandwidth, so simply filtering to separate a signal is no longer viable. Accommodating the full duplex portion of the spectrum involves the diplexer being replaced by a directional coupler. This directional couple can have at

most approximately 30dB of isolation between the upstream and the downstream direction in the full duplex band. The relatively low isolation in the directional coupler means that some of the downstream signal will leak into the upstream received signal. This leakage is especially problematic as the downstream signal has a much higher signal level than the upstream signal at the CMTS, which has been greatly attenuated in traversing the cable system from CM to CMTS. The echo paths and leakage through the hybrid are illustrated in Figure 1.5.

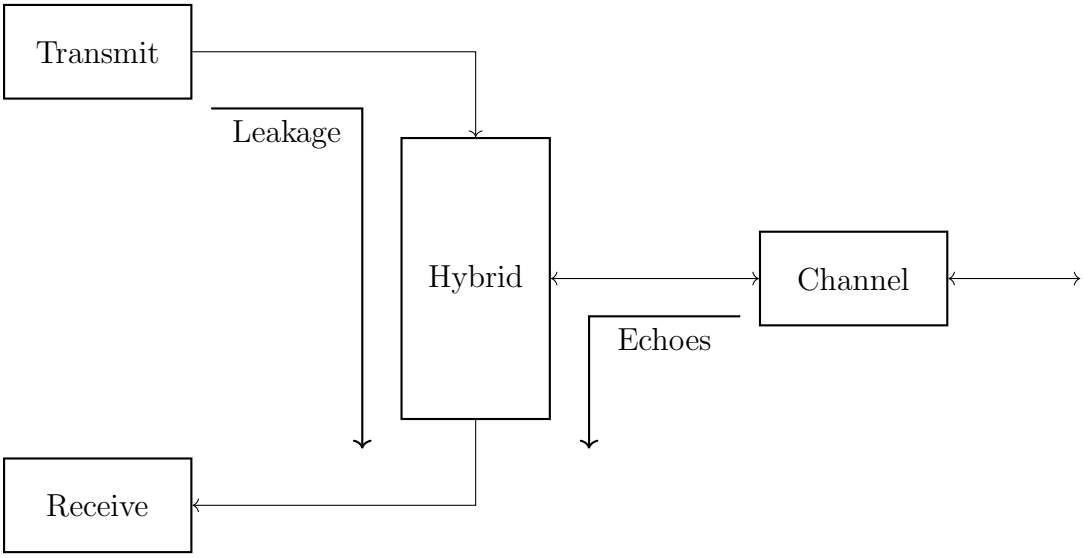

Figure 1.5 Echo Paths Inside the Node

The hybrid also has greater loss than the diplexer, which means that any power amplifiers in the processing chain must be driven higher. Driving the power amplifier higher causes it to become nonlinear by approaching its compression point, which adds more distortion into the higher amplitude portions of the signal. Due to this factor and a number of other noise sources on the upstream data such as noise from the PAs (Power Amplifiers), noise from the ADCs and DACs, and oscillator phase noise, the SNR on the upstream data is likely to be very low if not negative. A rough sketch of the spectrum is shown in Figure 1.6.

If the upstream signal is to be recovered the self interference from the downstream signal needs to be removed along with any echoes that the coaxial cable network introduces. It is for these reasons that an echo canceller is needed to restore the upstream signal. In

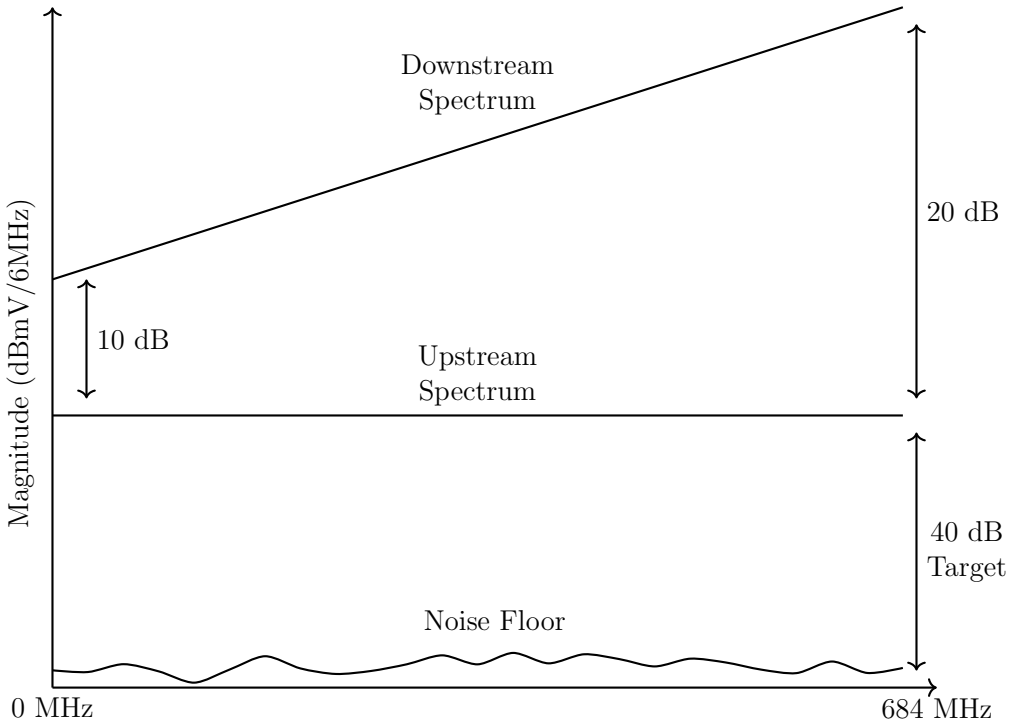

Figure 1.6 Received Spectrum

Figure 1.7, the location of the echo canceller in the node is shown. The function of the echo canceller is to exactly mimic the "echo transfer function" of the system so that the echo can be subtracted from the received signal. It is very difficult to measure the "echo transfer function" while receiving the upstream signal as the upstream signal will bury the echo. The echo canceller can also be thought of as a self interference canceller, as it also has the job of removing the portion of the transmitted downstream signal that "leaks" through the hybrid into the received upstream signal.

The echoes are introduced by the taps and CMs in the coaxial cable network, as well as any tiny imperfections present in the coaxial cables. As can be seen in Figure 1.8, the cable network is made up of lengths of coaxial cable connected to a series of taps (directional couplers), which take some of the signal and couple it to the attached cable modems (CM in the diagram), which would be located inside the users home. The coaxial cable itself is lossy, meaning that the signal will become increasingly attenuated as it moves through the network. To help make up for this loss, as well as for the power that is lost from the taps,

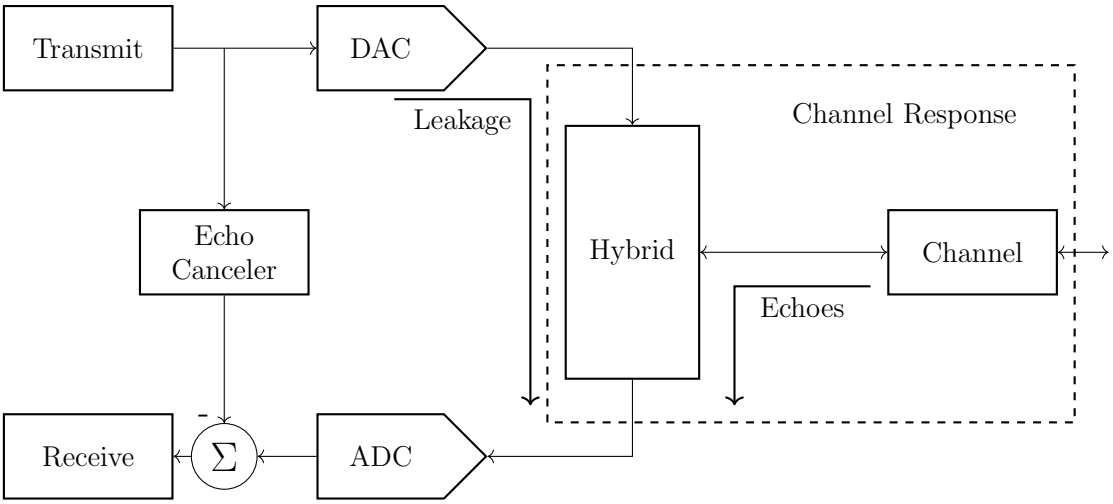

Figure 1.7 Node Diagram With Echo Canceller

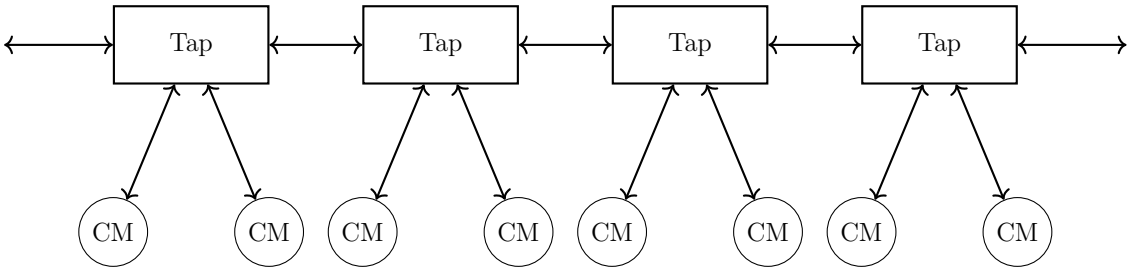

Figure 1.8 Coaxial Cable Network

the fraction of the power being diverted in each tap depends on the location of the tap in the network. The taps are arranged so that each cable modem gets the same amount of power. Unfortunately, these taps are not perfect by nature and generate reflections that show up at the CMTS as echoes. As the downstream signal hits each of the taps an impedance mismatch means that the tap cannot absorb all of the incoming power and thus some is reflected.

1.5 Statement of the Problem

The cable industry has planned for full duplex to be the next step in the evolution of DOCSIS. The DOCSIS 3.1 full duplex standard was released in October 2017, its feasibility was based entirely on theory and simulation and also under the perhaps shaky assumption that advanced hardware could be developed that supports it. This being the case, there are still a few unresolved issues around implementing a full duplex system in hardware. These

issues mostly center on the echo canceller. The echo canceller needs a significant time to determine the echo response of the channel. If the echo response of the channel changes with time as it is certain to do, the rate of change may be too fast to be considered constant over a defined "dwell time". This results in an improperly canceled signal and a unacceptably high noise interference. The main objective of this research is to determine the rate of change of the channel and the effect this rate of change has on the echo cancellation. To accomplish this a system will be designed and implemented on an FPGA that measures the characteristics of the channel as a function of time and determines the residual effects after echo cancellation.

In order to determine the characteristics of the channel that are most susceptible to rapid change, a model of a full duplex cable network is developed. Mathematical descriptions of the significant components in a full duplex cable network must be developed in order to properly model transmission in the system. These mathematical descriptions outline noise sources inherent to the system, so that their effect on echo cancellation can be accounted for. Once a detailed model of the system is developed it will be clear what signals the echo canceller needs to remove. The effect of the channel changing with time should be clear and the parameters associated with the change should be revealed. The measurement system can then be designed to find to these parameters.

1.6 Thesis Outline

The remainder of the thesis is organized as follows. Chapter 2 goes into detail on the full duplex DOCSIS network and provides mathematical description and analysis of the critical components that make it up. At the end of Chapter 2, analysis is done on the received upstream signal in the full duplex node to highlight the noise corrupting it. The result of this analysis is an equation showing the different impairments on the received signal. Chapter 3 uses this equation to show the need for self interference (echo) cancellation in the Full Duplex Node. Chapter 3 also provides analysis on echo cancellation and shows that it is limited by phase instability. The limitation that phase shifts imposes on echo cancellation is then

expanded into showing the need for a stable channel response for successful echo cancellation. Finally, Chapter 3 parameterizes the channel and provides a solution for characterizing the channel.

Chapter 4 analyzes a Phase Locked Loop based method selected to measure the channel parameters of importance outlined in 3. Chapter 5 goes through the issues of a high performance PLL circuit that is needed to measure the parameters of interest. Chapter 5 discusses the hardware implementation of the PLL as well as design procedures for selecting the parameters of the loop. Special care is given to the selection of the word lengths used in the PLL implementation on an FPGA.

The objective of Chapter 6 is to verify the design from Chapter 5. The chapter verifies that the Chapter 5 design meets all the goals outlined in Chapter 3 by simulating both the network and the PLL in MATLAB. Once the design is verified in MATLAB, it is built on an FPGA and further verified. The FPGA implementation shows that the design is able to meet the design goals. A series of measurements are explained, performed, and analyzed in Chapter 7. The measurements performed further show the success of the system as well as providing insight into the research goals of this thesis. Finally, the results of the research are summarized in Chapter 8. Chapter 8 concludes the work done in the thesis as well as outlining future work that can be performed in this area of research.

2. Critical Components of a Full Duplex Network

In this chapter the operation and limitations of a full duplex DOCSIS network is explored to provide the necessary background for the subsequent chapters. Chapter 3 will frame a clear statement of the problem using concepts explored in this chapter as a reference. This chapter begins by showing the architecture of the full duplex system specified in version 3.1 of DOCSIS. The system architecture and the mathematical models of its components are developed to give the research a firm theoretical basis. The mathematical descriptions of the components reveal some of the limitations of the full duplex DOCSIS network. After the relevant components are described the system is analyzed and the major issues explored.

2.1 Architecture of a Full Duplex Cable System

The architecture of a FDX DOCSIS network is shown in Figure 2.1. The node connects to and receives data from the internet. It maps the data into symbols and uses these mapped symbols to modulate a carrier. The down stream bound modulated carrier is transmitted over coaxial cable to Cable Modems (CMs) in subscriber households. In the opposite direction, the CMs in subscriber households generate data which modulates a carrier that is sent over the cable network to the CMTS. The network connects CMs in subscriber households to a node in the head end. The signals entering the network from the CMs to the node are considered to be in the upstream direction, and the signals sent from the node to network are considered to be in the downstream direction.

Inside the node section of Figure 2.1 there are a number of components. The data coming from the internet is first sent to the Media Access Control (MAC) layer of the downstream

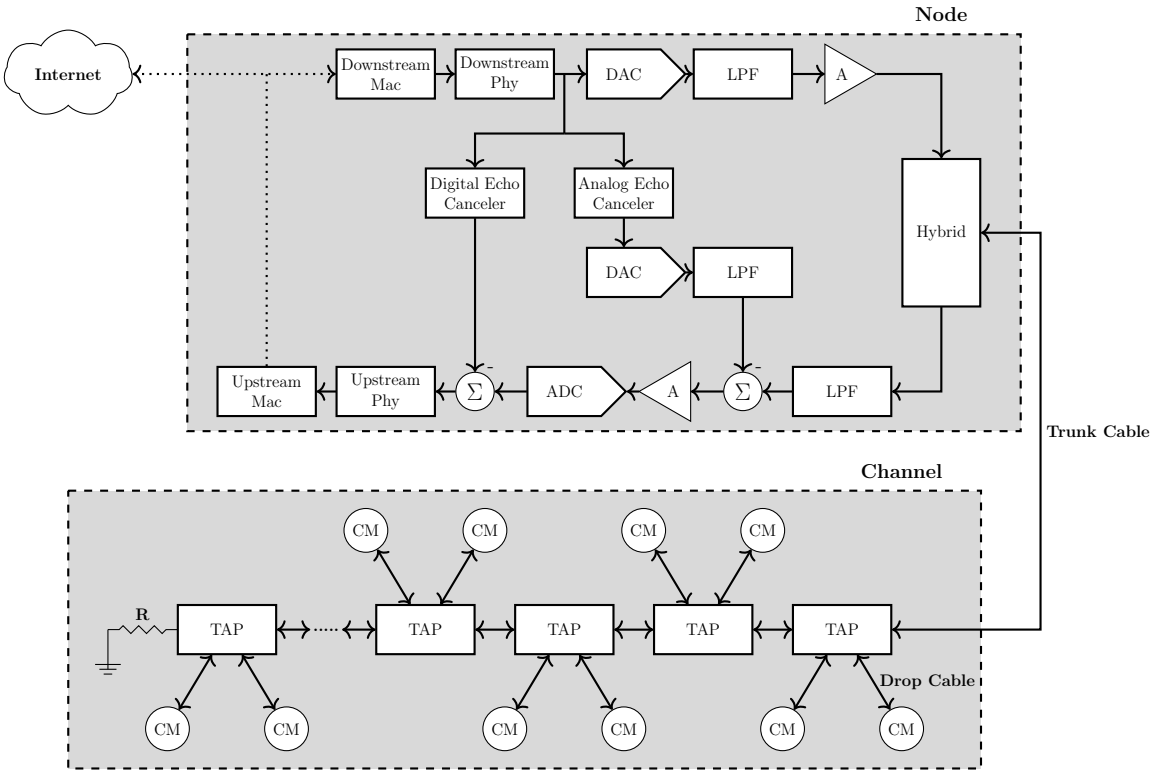

Figure 2.1 Full Duplex DOCSIS Network Architecture

path in the node. This MAC layer controls transmission of DOCSIS signals through a variety of algorithms which are summarized in Appendix 8.2. After the MAC layer, the data is passed into the Physical (PHY) layer. The function of this PHY layer is to convert the data into OFDM symbols. The PHY layer is also responsible for adding Forward Error Correction (FEC) to the incoming packets from the MAC layer [7]. After the data has been transformed into OFDM symbols, it is converted to an analog waveform in a Digital Analog Converter (DAC). This analog waveform is filtered to remove images and amplified to give it the power profile specified in the DOCSIS 3.1 standard. The signal is then sent through a directional coupler (Hybrid) and into the trunk cable.

In the upstream direction the processing in the node is in the reverse order. The upstream analog waveform is first band limited with a filter, passed through a low noise amplifier, and then converted to digital samples by an Analog to Digital Converter (ADC). The samples are filtered digitally and organized into OFDM symbols. These OFDM symbols are then

demodulated in the upstream PHY layer. The recovered data is then sent to the upstream MAC layer and from there onto the internet.

The transmitted signals are translated in frequency and placed into channels by various combinations of filtering, up conversion, and down conversion. A channel is the bandwidth allocated for one full duplex OFMDA spectrum. After the analog waveform has been organized into channels it is transmitted into a trunk cable. This cable then connects to a series of multi-port directional couplers, referred to as taps, with specific coupling strengths [9]. Cable modems located in the subscriber households are connected to the taps via drop cables. Figure 2.1 shows only two CMs attached to each tap, but the number of taps in a "tap" varies and depends on the density of subscribers. The end of the trunk is terminated to prevent impedance mismatches that generate interference in the form of echoes. Impedance mismatches will be discussed in detail later in this chapter.

2.2 Functional Operation of Critical System Components

This section of the chapter will describe the components mentioned in the previous section. Mathematical and functional descriptions of the components will shed light on the issues faced in the full duplex DOCSIS network. Only the blocks that directly relate to the measurements of interest done in this thesis are described, descriptions of the MAC layer, DOCSIS PHY layer components, taps, and cable modems can be found in Appendix 8.2.

For the purposes of this research, the transmit signal will be a single cosine tone at some arbitrary frequency within the FDX band. The analysis determined in the subsequent sections are done in the general case. The results of these analysis are later used to apply to the single test tone operation.

2.2.1 Downstream PHY Layer

The physical layer of the DOCSIS 3.1 node is mainly responsible for converting the incoming binary data into OFDM symbols. The physical layer also maps packets received

from the MAC layer into FEC codewords [7]. A simplified block diagram of the DOCSIS 3.1 PHY layer is illustrated in Figure 2.2. Figure 2.2 omits a few details of the actual processing path, mainly the PLC (Phy Link Layer), of the DOCSIS 3.1 PHY layer but the essentials for modulation are all present. Descriptions of the significant components shown in this diagram can be found in Appendix 8.2. Readers unfamiliar with DOCSIS cable networks may need to read this appendix to fully understand the remainder of this chapter.

The word "symbol" is widely used in the literature but it is used to refer to different things, before continuing it is of worth to define exactly what is meant by a "symbol" as it will be used in this thesis. From Figure 2.2 [7] it is seen that the IFFT block takes in parallel streams of N_{FFT} complex amplitudes generated by the constellation mapper. This group of N_{FFT} parallel complex amplitudes going into the IFFT block are collectively referred to as a frequency domain OFDM symbol. The output of the OFDM modulator after the parallel to serial converter is a set of $N_{FFT} + N_{CP} + N_{RP}$ complex samples. The core of which is the N_{FFT} samples generated by the IFFT block. A cyclic prefix of N_{CP} is prepended to the core as well as a role off period of N_{RP} samples. This entire set of $N_{FFT} + N_{CP} + N_{RP}$ complex samples are referred to as an OFDM symbol. Every frequency domain OFDM symbol is translated to an OFDM symbol that is sent to the frequency up converter.

2.2.2 Converting a Complex Baseband to a Real RF

A crucial step in the formation of the DOCSIS spectrum is not shown in Figure 2.2. The processing path shown in Figure 2.2 outputs complex baseband OFDM symbols, this means that there is a real and imaginary component (also referred to as the Inphase and Quadrature components) and the spectrum of the signal only occupies the band from 0 to 192 MHz. However, in Figure 1.3 the region from 5 to 1218 MHz is shown as occupied. The broader spectrum is filled by parking a number of different complex baseband OFDM spectrums in different segments of the overall spectrum using a process called frequency up conversion. Each full duplex DOCSIS node produces a number of complex baseband OFDM spectrums that are up converted to different bands. This conversion is accomplished with

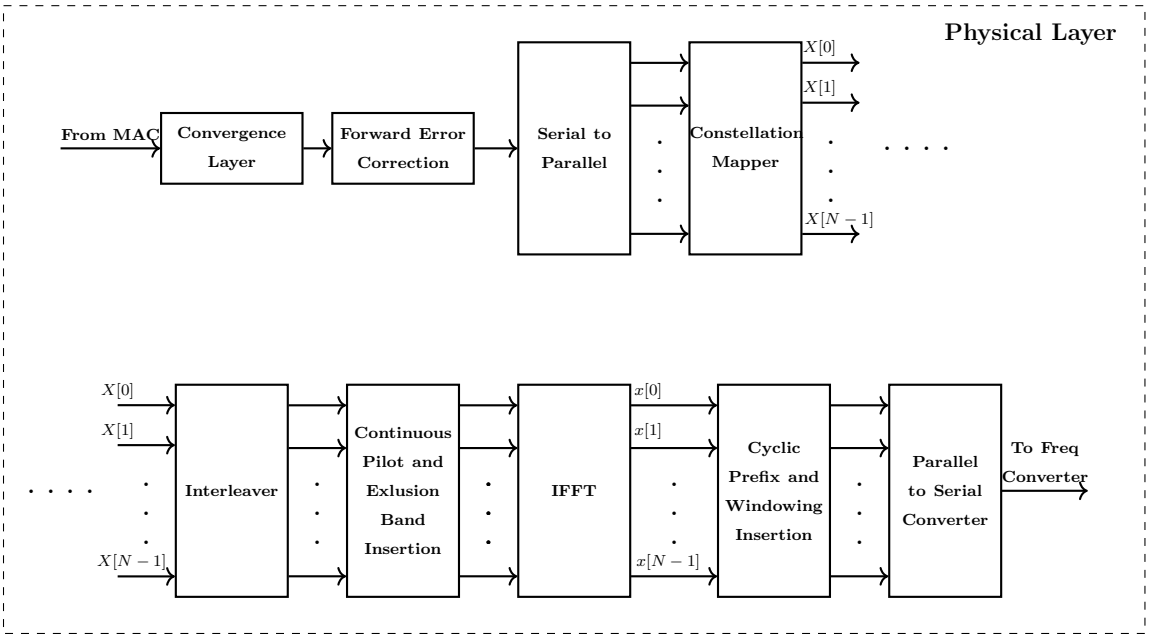

Figure 2.2 Simplified Downstream PHY Processing

the circuit shown in Figure 2.3. A frequency up converter separates the complex baseband signals into real and imaginary parts, multiplies them by quadrature sinusoids, then sums them to get a real bandpass signal.

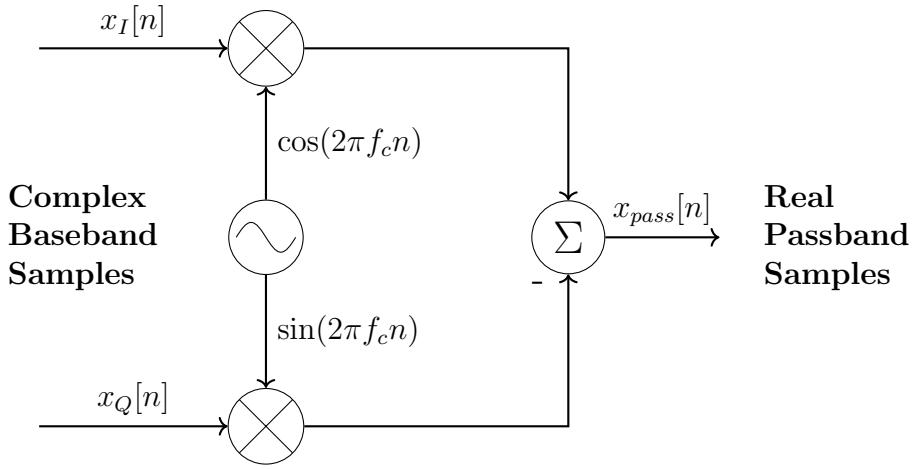

Figure 2.3 Frequency Up Converter

The output of the frequency up converter as shown in Figure 2.3 is given by the following equations.

$$\begin{aligned} x_{pass}[n] &= Re\{(x_I[n] + jx_Q[n])e^{j2\pi f_c n}\} \\ &= x_I[n]\cos(2\pi f_c n) - x_Q[n]\sin(2\pi f_c n) \end{aligned} \tag{2.1}$$

The output of the frequency up converter is a real version of the input with its frequency now centered at F_c Hz.

The local oscillator in (2.1) is shown to output pure noise free sinusoids. In practice the local oscillator will be corrupted in both amplitude and phase. The output of the quadrature local oscillator will take the form [10]:

$$\alpha_{TX}[n] = (1 + a_{TX}[n])e^{(j2\pi f_c n + \phi_{TX}[n])}, \tag{2.2}$$

where $\phi_{TX}[n]$ is the phase noise of the carrier used to mix the incoming signal and $a_{TX}[n]$ is the amplitude noise. The phase noise in the local oscillator, which is constructed digitally using a circuit known as an NCO (Numerically Controlled Oscillator), is due to truncation in the phase accumulator and the amplitude noise is due to truncation in the output tone. It is clear from (2.2) and (2.1) that frequency up conversion introduces phase noise to the signal and is therefore an imperfect process. It is important that this oscillator has as little phase noise as possible, as phase noise introduces inter carrier interference which will decrease the MER of the signal. More phase noise is introduced in the frequency down conversion process. The down converted signal takes on the following form:

$$x_{downconverted}[n] = x_{pass}[n] * h[n](1 + a_{RX}[n]])e^{(-j2\pi f_c n + \phi_{RX}[n])} \tag{2.3}$$

where $h[n]$, $\phi_{RX}[n]$, and $a_{RX}[n]]$ are the network response, the receiver oscillator phase noise, and the receiver oscillator amplitude noise respectively. Since the receive and transmit oscillators are locked to each other, $Re\{x_{downconverted}\}$ becomes the inphase (real) received symbols and $Im\{x_{downconverted}\}$ the quadrature (imaginary) received symbols. The phase noise terms introduced in the frequency down conversion will also contribute to inter carrier

interference and degrade the MER of the signal. This analysis shows that a significant amount of noise can be introduced in the frequency up conversion step of the DOCSIS processing chain.

2.2.3 Digital to Analog Conversion

A crucial component in the DOCSIS processing path is the Digital to Analog Converter (DAC). A DAC converts a digital signal to an analog signal, but the process is unfortunately not without error. There are two types of error introduced by a DAC: amplitude error and phase error. This subsection explains the conversion process from digital to analog in sufficient detail to provide a footing for the sources of noise introduced by the DAC and the effect that they have on the echo cancellation process.

Principle of Operation

A DAC can be thought of as having a sample and hold like function, although this is not truly the case. A DAC converts a value held in a digital register to a analog voltage level and holds that analog voltage level until the value in the digital register is updated. The process can be viewed as the DAC sampling the digital value and holding it until the digital value is updated, normally over a time period equal to the sampling period. In doing this, a DAC takes a discrete time sequence and converts it to an analog signal.

The resolution of the signal that is generated from the DAC depends on the number of bits used in the digital representation of the signal of interest and the sampling time, T_s, of the DAC. The more bits and the shorter the sample time the closer the output of the DAC will resemble the equivalent analog waveform. For example, a DAC that can output voltages between $\pm 5V$ and has a 3-bit register could only output 8 different voltages between $\pm 5V$. A DAC with a 14-bit register could output $2^{14} = 16384$ different voltages between $\pm 5V$. Having a finite number of output levels gives rise to quantization noise, which is discussed in further detail in Chapter 5. Figure 2.4 shows the DAC output (red) for an ideal signal (black) that was sampled at time nT_s and sent to the DAC.

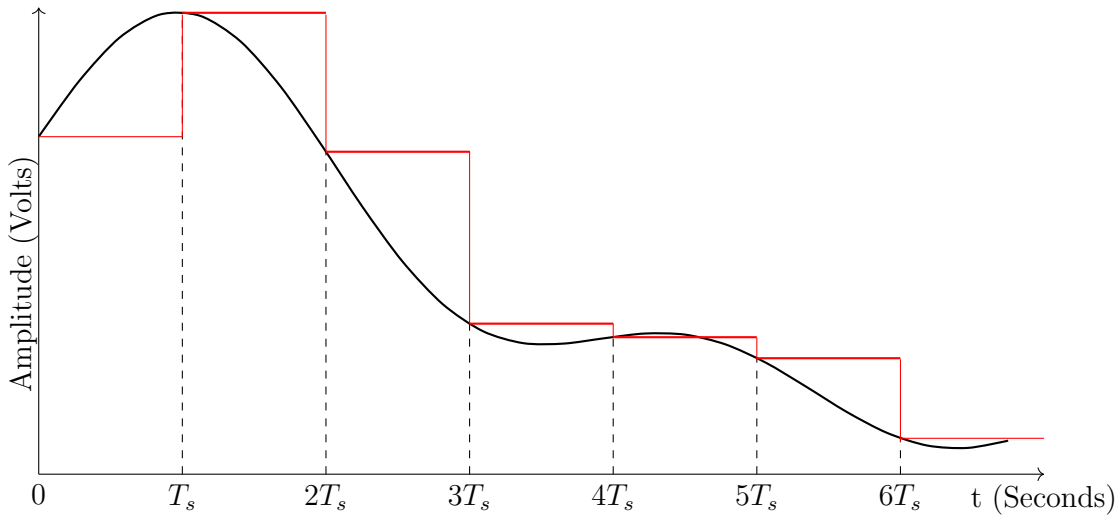

Figure 2.4 Input and Output for a DAC

The sample and hold like function of the DAC can be approximated with a purely theoretical circuit that operates as if the the digital signal was multiplied by a series of impulses centered at the sampling times $0, T_s, 2T_s, 3T_s, ...$ and then passed through a filter with an impulse response that is constant for one sample time and zero otherwise. In mathematical terms the impulse response of this constant response filter is defined as:

$$h(t) = \begin{cases} 1,\ 0 \leq t \leq T_s \\ 0, \text{Otherwise} \end{cases} \tag{2.4}$$

The frequency response of this filter is obtained by taking the Fourier transform of $h(t)$ and is given by

$$H(\Omega) = T_s e^{-j\Omega\frac{T_s}{2}} \left(\frac{\sin\left(\frac{\Omega T_s}{2}\right)}{\frac{\Omega T_s}{2}}\right) \tag{2.5}$$

where:

T_s = Sampling Time of the DAC in seconds

Ω = Analog Frequency in radians/second, $= 2\pi F$

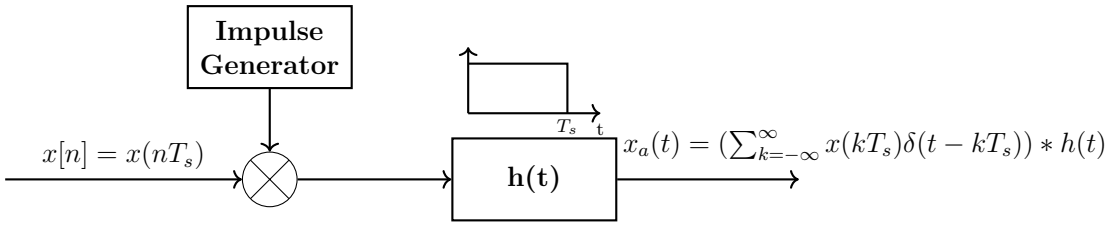

Figure 2.5 Equivalent System of Sample and Hold Dac

The equivalent system of the DAC is shown in Figure 2.5. The effect of impulse sampling places an infinite number of spectral copies of the original signal at integer multiples of the sampling rate. Therefore, the output of the impulse sampler in the frequency domain is the infinite sum of spectral copies, $\sum F_s X(\Omega - k\Omega)$. The filter with the impulse sampler reduces, but does not annihilate, these spectral copies. The spectral copies are removed by a low pass filter referred to as a reconstruction filter that is placed after the DAC. This reconstruction filter is shown as LPF in Figure 2.1 on page 14. The output of the system in Figure 2.5, which is the output of the DAC, has Fourier Transform

$$X_a(\Omega) = \sum_{k=-\infty}^{\infty} e^{\frac{-j\Omega T_s}{2}} \frac{\sin(\frac{\Omega T_s}{2})}{\frac{\Omega T_s}{2}} X(\Omega - 2\pi k F_s) \tag{2.6}$$

If the quantization noise of the signal is so fine it can be ignored and the reconstruction filter is perfect, which will never happen in practice, the output of the reconstruction filter is equal to the digital waveform sampled by the DAC. In which case the Fourier Transform of the output is

$$X_a(\Omega) = e^{\frac{-j\Omega T_s}{2}} \frac{\sin(\frac{\Omega T_s}{2})}{\frac{\Omega T_s}{2}} X(\Omega) \tag{2.7}$$

Clock Jitter Noise

The analysis for the output of the DAC was done for an ideal clock source where the sampling time is always at exact integer multiples of T_s. In reality the clock will have some jitter associated with it. This jitter in the edges of the clock introduces distortion to the

output of the DAC that cannot be removed. The time deviation between sample time nT_s and the corresponding rising edge of the clock is symbolized

$$\Delta\mathbf{t}[n] = \frac{\Delta\boldsymbol{\phi}[n]}{2\pi F_0} \tag{2.8}$$

where:

$\Delta\mathbf{t}[n]$ = Sequence of Random Variables in Units Seconds
n = Sample Number
$\Delta\boldsymbol{\phi}[n]$ = Phase Deviation in Radians at Sample n
F_0 = Clock Rate in Hertz

The phase deviation of the clock is usually specified in terms of its rms jitter in seconds, i.e in terms of $\sqrt{\overline{\Delta\mathbf{t}^2[n]}}$. This jitter causes the time of the positive edge of the clock to change, which changes the length of the "hold" of the sample and hold of the DAC. This is made clear in Figure 2.6, where the jitter delays the clock edge that clocks $x[n]$ into the register. In the case of Figure 2.6, the jitter causes $x[n-1]$ to remain in the register by $\Delta\mathbf{t}[n]$ too long.

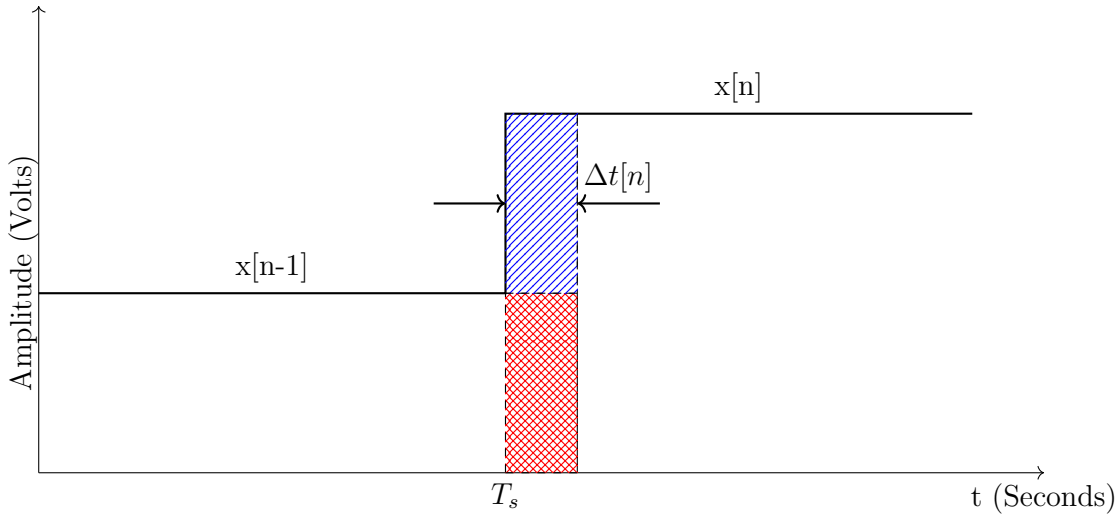

Figure 2.6 Sampling Distortion Caused by Timing Jitter

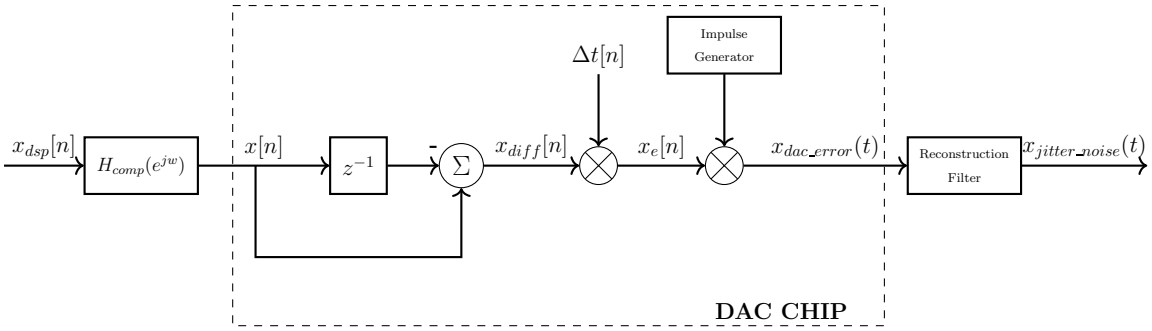

Figure 2.7 Clock Jitter to Additive Noise Model

The error at sample time nT_s is illustrated in Figure 2.6 by the rectangle of width $\Delta t[n]$ and height $x[n] - x[n-1]$ that is shaded with blue diagonal hatching. The practical DAC output is that of the ideal DAC output less the error, i.e. minus the diagonal blue hatched pulse.

Since $\Delta \mathbf{t}[n]$ is certain to be much smaller than T_s, the blue diagonal hatched pulse can be modeled as a delta function with area $(x[n]-x[n-1])\Delta t[n]$ and located at $t = nT_s$. That is, the error introduced by the jitter at nT_s can be modelled by $(x[n]-x[n-1])\Delta t(nT_s)\delta(t-nT_S)$, that being the case, the error sequence is given by the following sum

$$x_{dac\_error}(nT_s) = \sum_{n=-\infty}^{\infty} (x[n] - x[n-1])\Delta t(nT_s)\delta(t - nT_s) \tag{2.9}$$

The model that converts clock jitter, $\Delta \mathbf{t}[n]$, into additive noise is presented as a block diagram in Figure 2.7. The circuit enclosed in the dashed rectangle is an implementation of (2.9), which is a model for the DAC with $x[n]$ being the digital input to the DAC and $x_{dac\_error}(t)$ being the additive noise corrupting the output of the DAC. The block before and the block after the DAC are an integral part of the digital to analog conversion, but are not built on the actual DAC chip.

The block that precedes the "DAC CHIP" is a compensation filter. The purpose of the compensation filter is to correct for the in-band distortion that the DAC introduces to the signal of interest. The in-band distortion appears as a $\frac{\sin(\pi F/F_s)}{\pi F/F_s}$ roll off, which attenuates

higher frequencies. The compensation filter that corrects for this roll off is a linear phase filter with frequency response

$$H_{comp}(e^{jw}) = \begin{cases} \frac{2\pi F/F_s}{1-e^{j2\pi F/F_s}};\ 0 \leq F \leq \frac{F_s}{2} \\ 0; \text{Otherwise} \end{cases} \tag{2.10}$$

The magnitude response of this filter is $\frac{\pi F/F_s}{\sin(\pi F/F_s)}$, which is the multiplicative inverse of the DAC roll off.

It is of interest to determine the power spectral density of the jitter noise $x_{jitter\_noise}(t)$, which is the noise caused by jitter that will be added to the output of the DAC. Since the reconstruction filter is a low pass filter with a flat pass band its output has the same power spectral density as its input over its pass band and is 0 for frequencies in the stop band, assuming an ideal filter. Figure 2.7 shows the input to the reconstruction filter to be $x_{dac\_error}(t)$. Therefore the PSD of $x_{jitter\_noise}(t)$ is that of $x_{dac\_error}(t)$ over the band 0 to $F_s/2$ and 0 otherwise.

Figure 2.7 shows that $x_{dac\_error}(t)$ is a train of impulses multiplied by the digital signal $x_e[n]$. Since multiplication in the time domain is convolution in the frequency domain and the Fourier Transform of the train of time domain impulse is the set of frequency domain impulses $\frac{1}{T_s}\delta(F - nF_s)$ for $n = ..., -1, 0, 1, ...$, the PSD of $x_{dac\_error}(t)$ is an infinite set of frequency shifted versions of the PSD of $x_e[n]$ with each being scaled by $\frac{1}{T_s^2}$. Therefore, the PSD of $x_{jitter\_noise}(t)$ is that of $x_e[n]$ scaled by $\frac{1}{T_s^2}$.

$$S_{x\_jitterx\_jitter}(2\pi F) = \frac{S_{xexe}(e^{j2\pi F/F_s})}{T_s^2} \tag{2.11}$$

The problem of finding the PSD for $x_{jitter\_noise}(t)$ is reduced to finding the PSD of $x_e[n]$. Continuing to work back in Figure 2.7 has $x_e[n] = x_{diff}[n]\Delta t[n]$. Under conditions where $\Delta t[n]$ is independent of $x_{diff}[n]$, which it always will be, it can be shown that the power spectral density of $x_e[n]$ is the convolution of the PSDs of $x_{diff}[n]$ and $\Delta \mathbf{t}[n]$. From this we

get the following equation for the power spectral density

$$S_{xexe}(2\pi f) = \int_{-\infty}^{\infty} S_{xdiffxdiff}(e^{j(2\pi f - 2\pi\lambda)}) S_{\Delta t\Delta t}(e^{j2\pi\lambda}) d\lambda \tag{2.12}$$

where $f = \frac{F}{F_s}$ and λ have units of cycles/sample, $S_{xexe}(e^{j2\pi f})$ is the PSD of $x_e[n]$, $S_{xdiffxdiff}(e^{j2\pi f})$ is the PSD of $x_{diff}[n]$, and $S_{\Delta t\Delta t}(e^{j2\pi f})$ is the PSD of $\mathbf{\Delta t}[n]$.

Continuing to work back to $x_{dsp}[n]$ shows $x_{diff}[n]$ is linked to $x_{dsp}[n]$ by the system function $H(z) = H_{comp}(z)(1 - z^{-1})$

$$\begin{aligned} H_{diff}(e^{j2\pi f}) &= H_{comp}(e^{j2\pi f})(1 - e^{-j2\pi f}) \\ &= \frac{2\pi f}{1 - e^{-j2\pi f}}(1 - e^{-j2\pi f}) \\ &= 2\pi f \end{aligned} \tag{2.13}$$

Since, $S_{xdiffxdiff}(e^{j2\pi f}) = |X_{diff}(e^{j2\pi f})|^2$,

$$\begin{aligned} S_{xdiffxdiff}(e^{j2\pi f}) &= |H_{diff}(e^{j2\pi f})|^2 |X_{dsp}(e^{j2\pi f})|^2 \\ &= |H_{diff}(e^{j2\pi f})|^2 S_{xdspxdsp}(e^{j2\pi f}) \\ &= 4\pi^2 f^2 S_{xdspxdsp}(e^{j2\pi f}) \end{aligned} \tag{2.14}$$

An expression for the PSD of $x_{jitter\_noise}(t)$ is found by a series of back substitutions. First $S_{xdspxdsp}(e^{j2\pi f})$ is used to get $S_{xdiffxdiff}(e^{j2\pi f})$ and then $S_{xdiffxdiff}(e^{j2\pi f})$ is back substituted to get the PSD for $x_{jitter\_noise}(t)$. After the chain of substitutions are made the equation becomes,

$$S_{xjitterxjitter}(e^{j2\pi f}) = \frac{1}{T_s^2} \int_0^{0.5} (2\pi\lambda)^2 S_{xdspxdsp}(e^{j2\pi\lambda}) S_{\Delta t\Delta t}(e^{j2\pi f - 2\pi\lambda}) d\lambda \tag{2.15}$$

A special case of great importance to this thesis is $x_{dsp}[n] = \cos(2\pi f_0 n) = \cos(2\pi F_0/F_s n)$. In this case $S_{xdspxdsp}(2\pi F) = \frac{1}{2}\delta(F/F_s - F_0/F_s)$ or $\pi\delta(2\pi F/F_s - 2\pi F_0/F_s)$ and $S_{xjitterxjitter}(2\pi F) =$

$\frac{(2\pi F_0/F_s)^2}{2} S_{\Delta t \Delta t}(e^{j2\pi(F-F_0)/F_s})$, which is simply $S_{\Delta t \Delta t}(e^{j2\pi F/F_s})$ shifted in frequency by F_0 and scaled by $\frac{1}{2}(2\pi F_0/F_s)^2$.

The equation for the PSD of $x_{jitter\_noise}(t)$ can be verified through use of Parseval's Theorem [11], by using a time domain analysis to get the total power in $x_{jitter\_noise}(t)$. Since $H_{diff}(e^{j2\pi f}) = 2\pi f$, $x_e[n]$ has a particular relationship to $x_{dsp}[n]$. A system with a transfer function of $2\pi f$ is known as a differentiator. Therefore $x_e[n]$ can be obtained by sampling the analog signal $\frac{d}{dt} x_{dsp}(t)$, which is the derivative of the analog equivalent of $x_{dsp}[n]$. From this the following is obtained

$$x_{diff}[n] = x'_{dsp}(t)\big|_{t=nT_s} = x'_{dsp}[n] \tag{2.16}$$

Since $x_e[n] = x_{diff}[n]\Delta\mathbf{t}[n]$, the power in $x_e[n]$ is

$$\overline{x_e^2[n]} = \overline{x'_{dsp}[n]^2 \Delta\mathbf{t}[n]^2} \tag{2.17}$$

and since $\Delta\mathbf{t}[n]$ is independent noise

$$\overline{x_e^2[n]} = \overline{x'_{dsp}[n]^2} \cdot \overline{\Delta\mathbf{t}[n]^2} \tag{2.18}$$

expressing the powers of $x'_{dsp}[n]$ and $\Delta\mathbf{t}[n]$ in terms of their PSD has

$$\overline{x_e^2[n]} = \int_0^{1/2} (2\pi f)^2 S_{xdspxdsp}(e^{j2\pi f}) df \cdot \int_0^{1/2} S_{\Delta t \Delta t}(e^{j2\pi f}) df \tag{2.19}$$

In the special case where $S_{xdspxdsp}(e^{j2\pi f}) = \frac{1}{2}\delta(f - f_0)$,

$$\overline{x^2_{jitter\_noise}(t)} = \overline{x_e^2[n]} = \frac{1}{2}(2\pi f_0)^2 \int_0^{1/2} S_{\Delta t \Delta t}(e^{j(2\pi f - 2\pi f_0)}) df \tag{2.20}$$

Which is the same total power obtained before. The term $x_e[n]$ appears as additive noise on the output of the DAC such that the expression at the output is

$$x_a(t) = x_{dsp}[nT_s] + x_e[nT_s] \tag{2.21}$$

These results now characterize the noise that is introduced to the transmitted signal by clock jitter noise in the DAC. This clock jitter noise will always be present, and is a source of noise that the echo canceller cannot compensate for.

2.2.4 Analog to Digital Conversion

The full duplex node both transmits the downstream signal and receives the upstream signal. An Analog to Digital Converter (ADC) is required to receive the upstream signal and demodulate it back into digital packets. An ADC samples an analog voltage and converts it into a digital word. The more bits used in the digital representation of the sample, the less error in the quantized sample. An ADC will always introduce some quantization noise into the digital sample. The noise introduced by the ADC is assumed to be zero mean, white, and Gaussian, with a variance that characterizes the ADC's Effective Number of Bits (ENOB). The variance of the error, which is its average AC power, is found as follows

$$\sigma^2 = \frac{1}{3} \cdot \frac{1}{2}^{2 \cdot \text{ENOB}} \cdot (1 - \frac{1}{4}^{\text{ADC BITS} - \text{ENOB}}) \, V^2 \tag{2.22}$$

where:

σ^2 = Variance of Gaussian Noise and Quantization Noise Added by ADC
ADC BITS = ADC Resolution in Bits
ENOB = ADC Effective Resolution in Bits

The expression in (2.22) is found from the quantization error introduced from converting an infinite precision signal to a finite precision signal. The ENOB is essentially how many of the ADCs rated bits are actually usable information and how many are entirely consumed by the noise. The ENOB is normally given by the manufacturer of the ADC in the device's data sheet.

The dynamic range of an ADC is the range between the noise floor of the device and the maximum level it can output. The dynamic range is essentially the range of voltage levels that the ADC can resolve into digital words. An ADC with a dynamic range of 40dB can resolve voltages from x volts to 100x volts. The range of the ADC is important because in the full duplex DOCSIS network the range of signal strengths can vary greatly. The ADC receiving the upstream signal will likely need a very high dynamic range as the power leaking from the downstream through the hybrid will be very large and the upstream signal from the CMs very small.

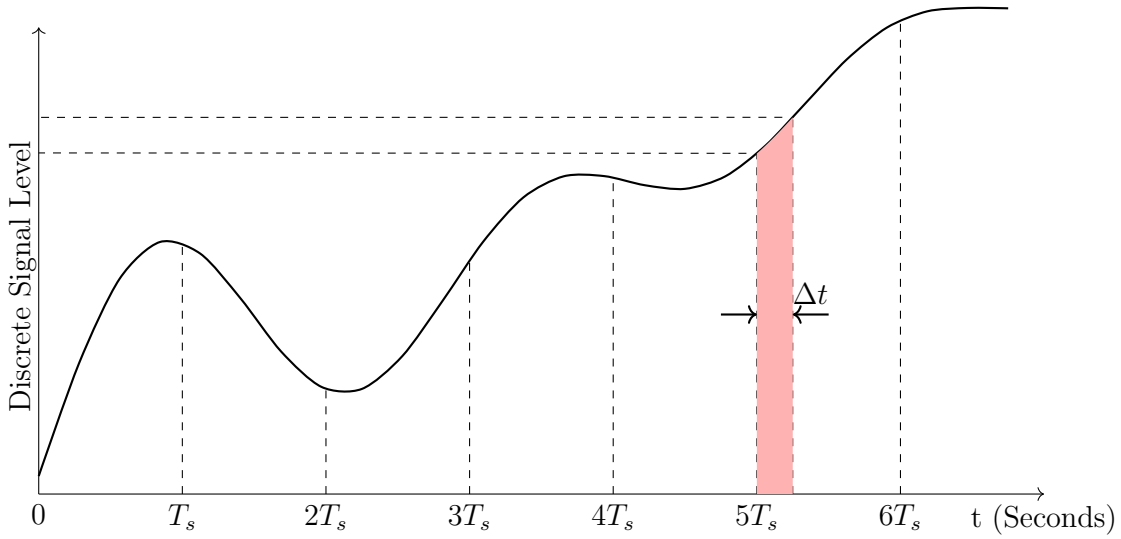

Figure 2.8 ADC Sampling Showing Error From Clock Jitter

Error Due to Clock Jitter

A second source of error is due to the error in sampling times. If the quantization noise is ignored the relation between the analog input of a ADC and the digital output is as follows

$$x[n] = x_a(t)|_{t=nT_s+\Delta t[n]} \tag{2.23}$$

where the $\Delta \mathbf{t}[n]$ is clock jitter, which was explained in the DAC section of this chapter. This

random clock jitter again causes error in our sampled signal. This error arises from irregular sampling as can be seen in Figure 2.8. If the error in sampling time, nT_s, is small the output of the ADC can be expressed with a first order Taylor series as follows

$$x[n] = x_a(t) + x_a^{'}(t)\Delta t[n]\Big|_{t=nT_s} \tag{2.24}$$

The error term in this signal is $x_a^{'}(t)\Delta t[n]\big|_{t=nT_s}$ which, surprisingly, is the exact same error term found for the DAC output given in Equation 2.21. This means that the error power in this signal is the same as that found in the DAC section in Equation 2.17 and thus further analysis is omitted.

If the same oscillator is used to clock both the ADC and the DAC, and the propagation time between the DAC and ADC is $<< T_s$, then only the DAC error affects the signal. If the propagation delay is $>> T_s$ or different clocks are used for the DAC and ADC the error will add on a power basis. If different oscillators are used for the ADC and the DAC then the error terms will be independent of each other and add on a power basis [10]. The reasoning for this is because the timing jitter for oscillators is due to internal thermal noise and imperfections in the crystals which will be unique to each crystal resulting in independent noise sources.

2.2.5 Hybrid Couplers

A major difference in the construction of an FDX node from the old HDX node is the replacement of a diplexer with a hybrid coupler. When the spectrum was divided as explained in Chapter 1, it was simple to separate the upstream signal from the downstream signal using low pass and high pass filters. The diplexer shown in Figure 2.9 cannot be used to isolate the upstream signals from the downstream signals when they both occupy the same bandwidth. A hybrid coupler is required to separate the upstream signals from the downstream signals when they occupy the same bandwidth.

A diplexer is a passive device containing a high pass and a low pass filter that combines two signals onto a common port. The high pass and low pass filters ensure that the signals

on the combined port do not interfere with each other as they will occupy different spectral regions. This was advantageous in the half-duplex system as the upstream was delegated to the lower frequencies and the downstream the higher frequencies, meaning that they could be transmitted and received on the same line and not interfere with each other. Typically, there is about 30dB of isolation between the input and output ports of the diplexer and 1.5dB insertion loss in the pass band frequencies [12].

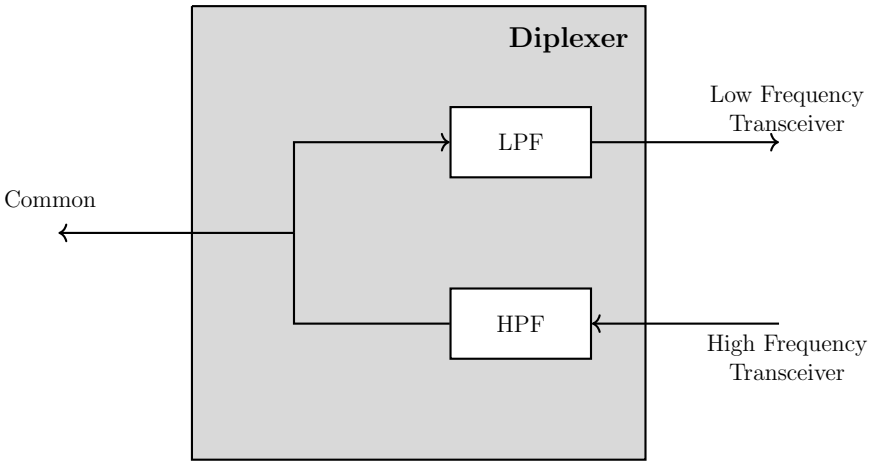

Figure 2.9 Diplexer Block Diagram

A hybrid coupler is actually just a 4-port directional coupler with one port terminated, and having the specific criteria of having a 3dB loss from each coupled port to the through port. The losses between ports is shown in Figure 2.10. For a directional coupler to be considered a hybrid coupler, the power entering the common port (2-wire side) must be split equally between the coupled ports (4-wire side ports). The issue in using a hybrid coupler arises from the imperfect isolation between the two coupled ports on the 4-wire side. This means that the transmit port will "leak" into the receive port. This leakage could be as little as 30 dB below the transmit power in a high quality device.

2.2.6 Coaxial Cable

The transmission medium used in the CATV network is coaxial cable. A coaxial cable is two concentric conductors separated by a solid insulator. The outer conductor is generally

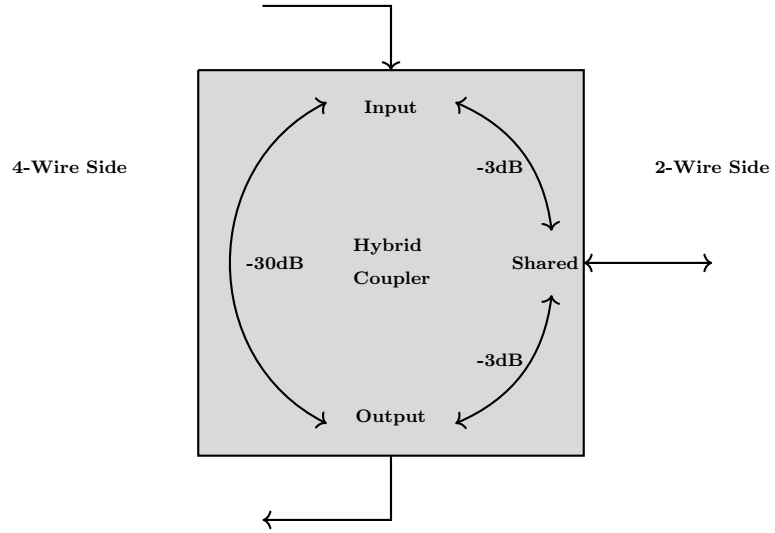

Figure 2.10 Hybrid Block Diagram Showing Typical Isolation and Insertion Loss

referred to as the sheath or shield and the inner conductor is known as the center wire. To minimize losses the solid insulator will have a very low dielectric constant and as such very low dielectric losses.

The discussion of the properties of the coaxial cable is important in the context of this thesis. The channel instabilities that are the subject of this thesis mostly arise from physical imperfections in the coaxial cable. Rapid heating, strong winds, or even cars driving over buried cable all have the possibility of introducing these channel instabilities, and as such it is important to understand the properties of the coaxial cables.

Attenuation

Generally, coaxial cable has a low pass type frequency response. There are different qualities of cable in the network as shown in Figure 2.1. Trunk cable has lower losses than drop cables, but is more expensive and considerably stiffer meaning it cannot be sharply bent without destroying the integrity of the cable. Trunk cable is used to connect the head end to the taps. The drop cables are cheaper, but are lower quality and have a bend radius of about 1.5 feet, so they are used for the relatively short connections from the taps to the CMs. The losses through a coaxial line are hard to completely characterize through a formula and are

generally analyzed through measured data. A frequency dependent attenuation constant, α, which has units nepers/meter approximates the attenuation in coaxial cables. It is roughly given by (2.25)

$$\alpha \approx \frac{\sqrt{LC}}{2}\left(\frac{R_{DC} + R'\sqrt{F}}{L} + \frac{G'F}{C}\right)\frac{nepers}{meter} \tag{2.25}$$

where:

α = Attenuation Constant in Nepers per Meter
L = Cable Inductance in Henrys per Meter
C = Cable Capacitance in Farads per Meter
R_{DC} = Cable DC Resistance in Ohms per Meter
R' = Cable AC Resistance in Ohms per Hertz
G' = Cable AC Conductance in Mhos per Hertz
F = Frequency in Hertz

The expression in (2.25) clearly shows that the loss, α, grows with frequency, showing the cable has a low pass type response.

The loss through a coaxial cable in dB is proportional to its length, meaning a 100 meter coaxial cable would have half as much loss as a 200 meter coaxial cable. It is for this reason that fibre has replaced the long stretches of trunk cable in the CATV plant. Cheaper drop cables are used to connect subscriber modem to the taps for reasons of economy. Each drop cable is relatively short, and it would appear that the quality of the cable is unimportant, but due to the shear number of drop cables used in a coaxial network being economical is extremely important.

To help to counteract the low pass effect of the coaxial cables the signals are transmitted with a positive tilt in the frequency domain with higher frequencies being transmitted at higher power. Such a tilt can be observed in Figure 1.6 on Page 9. The positive tilt means that the higher frequencies have more power directly out of the node and thus will have

greater interference on the received signal.

Characteristic Impedance

The characteristic impedance in coaxial cables is defined as the ratio of the complex amplitudes of the voltage of the center conductor with respect to the sheath over the amplitude of the current in the center conductor. DOCSIS CATV networks use cables with a characteristic impedance of 75 Ω [9]. The characteristic impedance of a cable is calculated using the following formula

$$Z_0 = \sqrt{\frac{R + j2\pi FL}{G + j2\pi FC}} \approx \sqrt{\frac{L}{C}}\,\Omega \tag{2.26}$$

where:

Z_0 = Characteristic Impedance in Ohms
L = Cable Inductance in Henrys
C = Cable Capacitance in Farads
F = Frequency in Hertz
R = Resistance in Ohms
G = Conductance in Mhos

Another important concept in coaxial cables is the reflection coefficient, Γ. Should the load not have an impedance equal to the characteristic impedance of the network, a percentage of the voltage going towards the load is reflected back towards the source. The reflections caused by this impedance mismatch are the source of the previously mentioned echoes. These echoes can reflect back and forth many times, leading to the possibility of many echoes on the received signal. The strength of these echoes relative to the forward traveling voltage is given by the reflection coefficient which is given by

$$\Gamma_L = \frac{Z_L - Z_0}{Z_L + Z_0} \tag{2.27}$$

where Z_L is the impedance of the load and Z_0 is the characteristic impedance. It can be seen from (2.27) that even a slight difference in impedance can lead to a non zero reflection coefficient. The reflection coefficient is generally complex and therefore has an amplitude and a phase component. Due to this the reflected and forward propagating waveforms are normally described by their amplitudes and phases. The complex amplitude of the voltage waveform in the "forward" direction is symbolized as $v^+(x_L)$ and the complex amplitude of the voltage waveform in the "reverse" direction is symbolized as $v^-(x_L)$, where x_L is the distance along the transmission line from the voltage source. In this case the relation between the complex amplitude of the forward and reverse travelling waveforms is as follows.

$$v^-(x_L) = \Gamma_L v^+(x_L) \tag{2.28}$$

Reverse propagating voltages can arise from even the tiniest imperfections in the load or even the cable itself. If the cable has some tiny flaw, such as an air bubble in the insulator between the center conductor and the shield, than a echo will be generated. These echoes (also known as micro-reflections) are the source of a large amount of distortion in DOCSIS networks and much work has gone into characterizing them [13].

Two port networks, such as the one in a DOCSIS network, are ubiquitously characterized by S parameters. The S parameters measure the gain through a 2-port device as well as its reflections. If a_1 is the complex amplitude of the incident wave in the forward direction, and a_2 is the complex amplitude of the incident wave in the reverse direction, then b_1 and b_2 are the complex amplitudes of the reflections, where b_1 is the reflection of the forward traveling wave and b_2 the reflection of the reverse traveling wave as illustrated in Figure 2.11. The complex amplitudes of the two reflected waves are given by

$$\begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \tag{2.29}$$

S-parameters for a 2-port network can be measured by a VNA (Vector Network Analyzer).

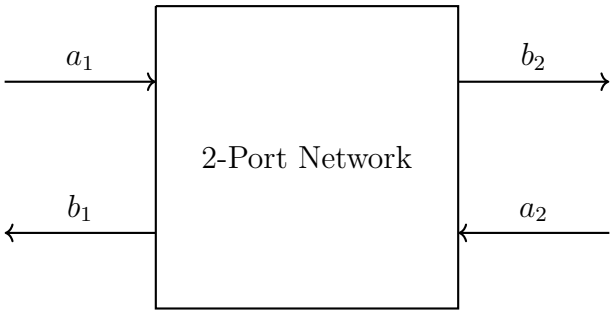

Figure 2.11 Generic 2-Port Network Showing S-parameter Usage

The goal of this research is to establish the amount of change in the channel the echo canceller can tolerate in a given time span. The S-parameters alone do not help achieve this goal as they are mostly a measure of forward and reflected gains. Another issue is that it is difficult to obtain the S-parameters accurately enough for the purposes of this research. The echoes are important to characterize for the CATV network, for which S-parameters are undoubtedly useful, however due to the issues outlined above the S-parameters will be largely ignored in this research.

Velocity of Propagation

Another important concept in coaxial cables is the velocity through which waveforms propagate through the network. The velocity at which the signal travels through the cable is shown as a percentage of the speed of light in a vaccuum. If V_p is 100% that means the signal is traveling at the speed of light, if it is 50% it is traveling at half the speed of light. The rate at which the signal propagates through the cable is dependent not only on the relative permeability and relative permittivity of the insulator used in the cable, but also on the straightness of the cable and the shield. A typical value for V_p of the coaxial used to connect to taps and CMs is 87% [7]. The formula for the velocity as a percent of the speed of light is given below.

$$V_p = \frac{100}{\sqrt{\mu_r \epsilon_r}} \text{ \% of Speed of Light} \tag{2.30}$$

where:

V_p = Velocity of Propagation as a Percent of the Speed of Light
μ_r = Real Part of the Relative Permeability of Dielectric Insulator
ϵ_r = Real Part of the Relative Permittivity of Dielectric Insulator

It is noted by taking only the real part of the dielectric constants μ and ϵ that the dielectric losses associated with the material are neglected. Ignoring these losses has very little effect on the outcome of the equation as they are normally very small in magnitude.

Physical forces that deform the cable even very slightly will change the velocity of propagation through the cable [14]. This is problematic as changes in the velocity of propagation changes the propagation delay of the waves in the cable. This change in propagation delay translates to phase shifts that manifest as noise in the signal. The echo canceller must adapt to these phase shifts, but they may happen too fast for the echo canceller to measure and react to them. It is of interest of this research to determine how consequential these phase shifts are on the overall performance of the system.

2.3 System Analysis of a Full Duplex Network

To demonstrate the need for an echo canceller an analysis of the full duplex network is required. This section will use the descriptions of the components in the full duplex DOCSIS network developed in the previous sections to derive a clear expression for the signal received at the full duplex node.

A simplified view of the full duplex system is shown in Figure 2.12. In Figure 2.12 all of the cable modems are clumped into one transmitting entity that transmits the signal $\widetilde{x_2}(t)$. The node is shown to contain both the transmit and receive portions. The input to the ADC, which is the output of the hybrid on the receive side of the full duplex node, is defined as $\widetilde{r}(t)$. The output of the DAC, which is the signal being transmitted from the full duplex node, is defined as $\widetilde{x_1}(t)$. To help distinguish the real pass band and complex base band signals, symbols for real pass band signals will have a tilde placed over them and

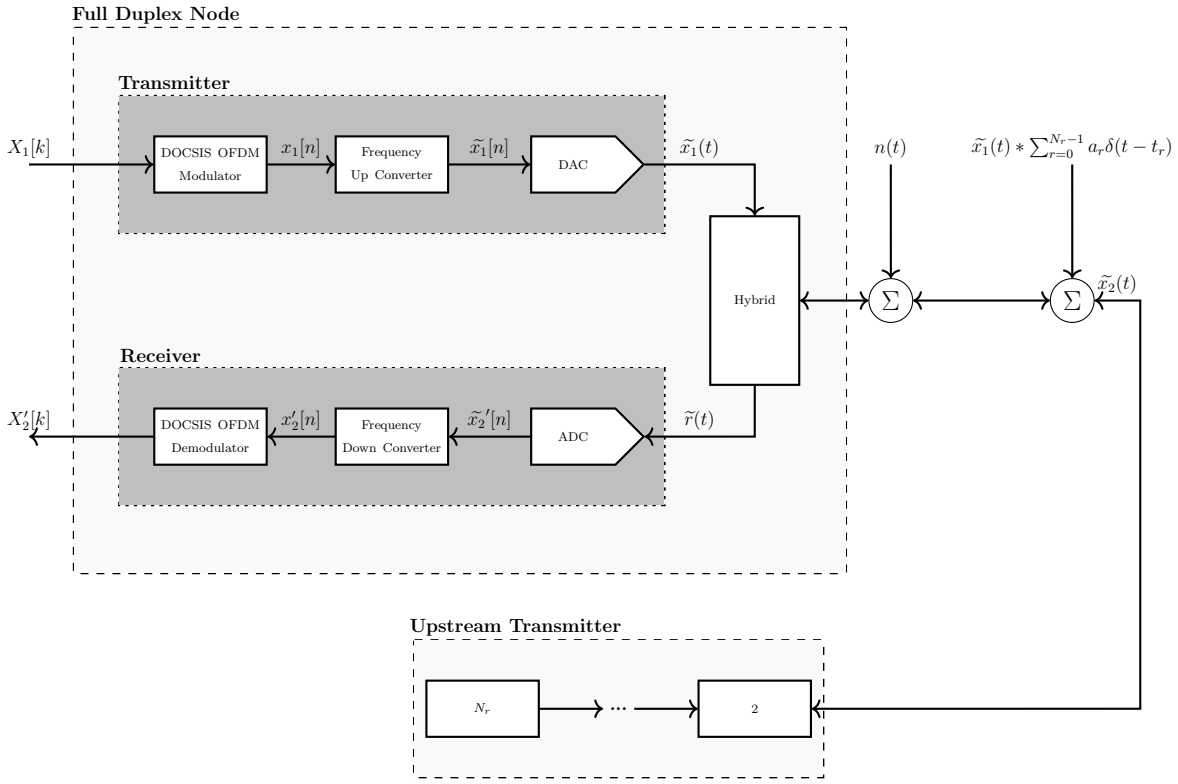

Figure 2.12 System Diagram of a Full Duplex Network

complex base band signals will not. With this notation it is clear that $\widetilde{x_1}(t)$ and $\widetilde{r}(t)$ are real pass band signals. The impulse response $h_{ij}(t)$ is the response of the channel between point i and j. The node itself is considered to be point 1 and the response $h_{11}(t)$ signifies the self interference channel in the node from the transmit to the receive side of the hybrid coupler. The response h_{21} characterizes the channel from the closest point of the upstream transmitter to the receive end of the hybrid coupler.

The additive noise $n(t)$ is assumed to be zero mean white Gaussian noise with known two sided power spectral density $N_0/2$ in units $\frac{\mathrm{W}}{\mathrm{Hz}}$. The second additive signal represents the reflections introduced by impedance mismatches in the channel. The simplified echo signal is given by the convolution $\widetilde{x_1}(t) * \sum_{r=0}^{N_r-1} a_r\delta(t-t_r)$, which is a sum of scaled and delayed versions of the transmitted signal. This represents the echoes which are effectively just scaled versions of the transmitted signal reflected back with a new amplitude and phase. The a_r

represents the amplitude of the echo of the transmitted signal. The delay term $\delta(t - t_r)$ characterizes the propagation time of the channel with t_r being the path delay. Although the reflection coefficient generally has an amplitude and phase associated with it, to simplify analysis a_r is a real scaling factor. The number of significant echoes in the channel is denoted N_r. In this analysis N_r is proportionate to the number of taps in the network.

The first step to be analyzed in the chain is the DOCSIS OFDM Modulator. The modulator takes in complex amplitudes $X_1[k]$, for sinusoids with frequencies $\frac{2\pi k}{N_{FFT}}$ and converts them into discrete time domain samples $x_1[n]$ for n $= 0, 1, ...N_{FFT} - 1$ which make up an OFDM symbol. The sequence $x_1[n]$ is given by the complex base band discrete time signal

$$x_1[n] = \frac{1}{\sqrt{N_{FFT}}} \sum_{k=0}^{N_{FFT}-1} X_1[k] e^{j\frac{2\pi n(k-N_{FFT}/2)}{N_{FFT}}},\ 0 \leq n \leq N_{FFT} - 1 \tag{2.31}$$

It is understood that a cyclic prefix and window are affixed to the discrete time domain symbols after this point, more information on this process can be found in Appendix 8.2. The signal is then frequency up converted in the digital domain to produce the real pass band signal $\widetilde{x_1}[n]$.

$$\widetilde{x_1}[n] = Re\{x_1[n]e^{-j2\pi F_c + \phi_T[n]}\} \tag{2.32}$$

The discrete time real pass band signal then is passed through a DAC to convert to continuous time. The output of the DAC, found using the formulae in the DAC section of this chapter, is given by

$$\widetilde{x_1}(t) = (\sum_{k=-\infty}^{\infty} \widetilde{x_1}[kT_s]\delta(t - kT)) * h_{DAC}(t) + \widetilde{x_1'}(t)\Delta\boldsymbol{t}(t) \tag{2.33}$$

where it is understood that $h_{DAC}(t)$ is the constant response filter explained in the DAC section of this chapter.

This signal then propagates through the hybrid, which is characterized by the response $h_{11}(t)$, and becomes a powerful self interference component in the received signal $\widetilde{r}(t)$. The self interference signal consists of the actual transmitted signal, along with noise terms that cannot be estimated in the receiver like the clock instability noise originating in the DAC and the phase noise originating from the oscillator. These components will degrade the SNR of $\widetilde{r}(t)$. The signal $\widetilde{x_1}(t)$ is transmitted at very high power relative to the upstream signal in $\widetilde{r}(t)$, but the attenuation through the hybrid coupler somewhat mitigates this. Even taking the coupler attenuation into account the self interference signal $\widetilde{x_1}(t) * h_{11}(t)$ will swamp the received signal.

Another major component in the received signal that must be taken into account is the echoes caused by impedance mismatches in the network. These echoes are not all significant and the most significant ones tend to be generated when the forward propagating waveform are reflected off of the taps. The transmitted signal $\widetilde{x_1}(t)$ hits these taps and a portion of it is reflected back to the node and corrupts the received signal further.

$$x_{echoes}(t) = \widetilde{x_1}(t) * \sum_{r=0}^{N_r-1} a_r \delta(t - t_r) \tag{2.34}$$

Each of the echoes has its own attenuation coefficient that is dependent on the location of the echo as well as a time delay that also depends on the location of the echo. These echos all manifest at the received signal $r(t)$ and are considered as removable sources of noise.

The upstream signal $x_2(t)$ is for the purposes of this analysis a summation of the transmitted signals from N_r cable modems. Each of the cable modems modulates and up converts the signals and sends them to a tap where they are all coupled together to generate the signal $\widetilde{x_2}(t)$. The channel response $h_{21}(t)$ is the overall response of this coupled signal and accounts for the individual responses from each cable modem to the full duplex node.

Taking all of the above into account it is possible to generate an expression for the signal at the receive port of the hybrid coupler. The expression is as follows

$$\widetilde{r}(t) = \widetilde{x_2}(t) * h_{21}(t) + \widetilde{x_1}(t) * \sum_{r=0}^{N_r-1} a_r \delta(t - t_r) + \widetilde{x_1}(t) * h_{11}(t) + n(t). \tag{2.35}$$

The received signal, $\widetilde{r}(t)$, consists of the desired signal, $x_2(t)$, frequency up converted and convolved with the channel response h_{21}, as well as being corrupted by the self interference signal, $\widetilde{x_1}(t) * h_{11}$, and a series of echoes of the transmitted signal. The received signal is also corrupted by zero mean wide band white noise $n(t)$. The signal $\widetilde{x_1}(t)$ itself is also corrupted by noise originating from imperfections in crystal oscillators that cannot be cancelled. This noise on $\widetilde{x_1}(t)$ will also appear in some form on the received signal $r(t)$. The channel response $h_{21}(t)$ is accounted for in adaptive equalizers in the full duplex node and adaptive pre equalizers in the cable modems. These equalizers diminish the distorting effect of the channel response.

The MER of the desired upstream signal $\widetilde{x_2}(t)$ must be high enough so that large QAM constellations are possible. The MER of a signal is also directly related to its SNR (Signal to Noise Ratio). If the upstream signal is received at a much higher power than the signals corrupting it then there is no issue and the receiver portion of the node will correctly recover the data. This is very unlikely to happen though as $\widetilde{x_2}(t)$ will be greatly attenuated as it propagates through the network. The self interference signal $\widetilde{x_1}(t) * h_{11}(t)$ is also highly problematic as $\widetilde{x_1}$ is transmitted at a very high power relative to the $\widetilde{x_2}$. $\widetilde{x_1}(t)$ also has a positive tilt applied to its spectrum to compensate for the low pass response of the channel, this means that in the higher frequency ranges the desired signal $\widetilde{x_2}(t)$ will likely have a negative SNR making successful recovery unachievable.

The final niggle in the received signal is the thermal noise $n(t)$. Thermal noise can be found in all communication systems and originates from random electron movement induced by heat [15]. Thermal noise is generally modeled as white which implies that its power spectral density is flat across the band of interest. In this case the band of interest is the full duplex band from 108 MHz to 684 MHz and it is reasonable to assume the thermal noise is flat over that range. Thermal noise over a limited bandwidth has a zero mean Gaussian distribution with a variance of σ_t^2.

Table 2.1 Minimum CNR for DOCSIS QAM Constellations

Constellation	**CNR (dB)**
QPSK	11.0
8-QAM	14.0
16-QAM	17.0
32-QAM	20.0
64-QAM	23.0
128-QAM	26.0
256-QAM	29.0
512-QAM	32.5
1024-QAM	35.5
2048-QAM	39.0
4096-QAM	43.0

The variance describes the AC power of the thermal noise. The value of the variance is determined by the CNR (Carrier Noise Ratio). The DOCSIS standard defines the minimum CNR allowed for different QAM constellations as listed in Table 2.1 [7]. The DOCSIS standard defines the CNR "as the ratio of the average signal power in the occupied bandwidth, to the average noise power in the occupied bandwidth given by the noise power spectral density integrated over the same occupied bandwidth" [7]. The variance of this noise in V^2 is given by

$$\sigma_t^2 = 10^{\frac{-\text{CNR}}{10}} \tag{2.36}$$

Now that the exact received signal is known and all of the major sources of noise identified, the concept of an echo canceller can be introduced in Chapter 3 so that the upstream signal can be recovered from $\widetilde{r}(t)$.

3. Self Interference Suppresion and Phase Instability

This chapter will examine the concept of self interference suppression and the effect that phase response instability has on it. The equations derived at the end of Chapter 2 show the received signal in a full duplex system is corrupted by a self interference signal, echo signals, clock noise, quantization noise, and AWGN. In this chapter, a method for removing the self interference signal and the echo signals will be examined. The limitations on the removal of these interference terms will be made clear. In particular the effect that variations in the phase response of the channel has on the self interference suppression will be shown.

3.1 Concept and Roll of Self Interference Suppression

This section will explain the algorithm behind self interference suppression. Chapter 2 has shown that some way of recovering the upstream signal from the noise introduced by the downstream signal must be developed. This is typically done with a self interference suppressor, which is also known as echo canceller. In the most basic terms, an echo canceller works by taking in a signal that is also propagating through a channel with a defined frequency response. The echo canceller attempts to emulate the response of the channel so that the difference between the output of the echo canceller and the echo generated by the channel is negligible. The basic algorithm for this process is shown in Figure 3.1 [16]. The structure used to emulate the channel response is known as an adaptive filter.

There are many different adaptive filter algorithms, but they generally follow 3 governing

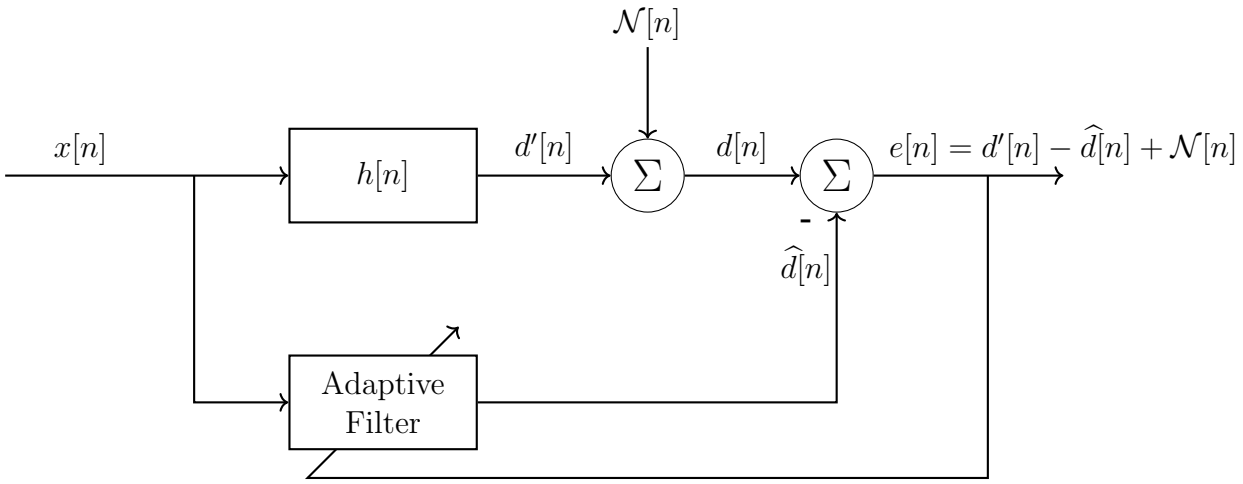

Figure 3.1 Adaptive Filtering

equations [17]:

$$\begin{aligned} y[n] &= w^T[n-1]x^*[n] \\ e[n] &= d[n] - y[n] \\ w[n] &= w[n-1] + f(x[n], e[n], \mu), \end{aligned} \tag{3.1}$$

where $w[n]$ is the taps or filter coefficients inside the FIR adaptive filter, μ is the update coefficient, and $f(*)$ is the function used for updating the coefficients. Different adaptive filter algorithms use different $f(*)$ with varying strengths and weaknesses. The most common and simplest type of adaptive filtering algorithms are LMS (Least Mean Squared). LMS adaptive filter adjust their coefficients in an attempt to minimize $\overline{e[n]^2}$ [17]. The update coefficient, μ, controls the step size of the adaptive algorithm with a larger μ meaning faster, but more coarse (higher steady state error) adaption. A smaller μ leads to slower, but finer (smaller steady state error) adaption. An example structure for the implementation of a FIR adaptive filter is shown in Figure 3.2, in this figure a bold type face is used to depict a matrix variable.

All adaptive filters require some time to train to the system that it is emulating. In DOCSIS the echo cancellers are given training sequences so that the adaptive filters can converge before transmission of data begins [7]. The time that the adaptive filters take to

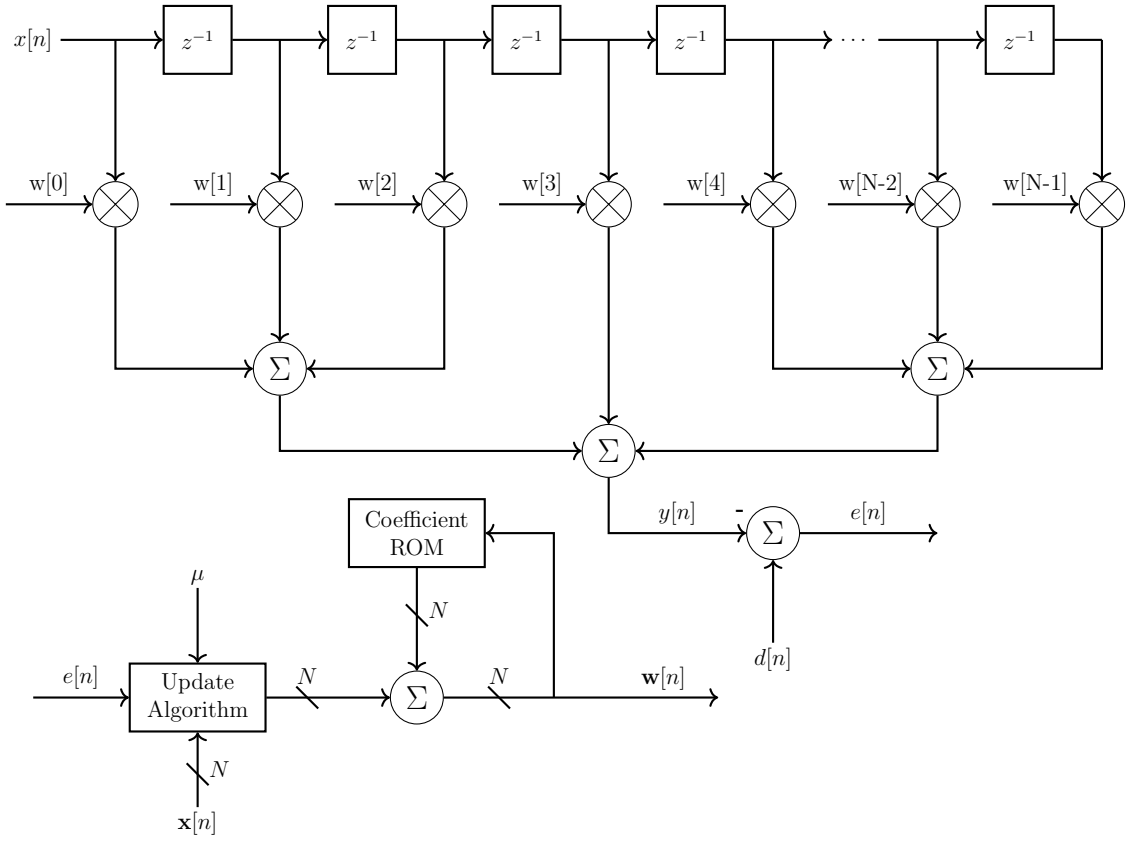

Figure 3.2 Adaptive Filter Structure

train to the system depends on the algorithm being used and the update parameter μ. The advantage of using an adaptive filter in the echo canceller is that after the filter has been trained to the system of interest it can further adapt if the system itself ever changes. This ability to re-adapt to new conditions is advantageous in the DOCSIS cable network as the cable network will change over time.

Adaptive filtering is the perfect solution to the self interference and echo interference outlined in Chapter 2. The filter can emulate the echo paths as well as the leakage paths through the hybrid. These interference components can then be removed from the received signal $\widetilde{r}(t)$ shown in Figure 2.12 on Page 37. With these interference components removed the recovery of $\widetilde{x}_2(t)$ is greatly improved.

Suppose that the response to an impulse applied to the transmit side of the hybrid coupler

in Figure 2.12 at the receiver side is given by $h(t)$. This response would include the echo and leakage paths seen by $\widetilde{x}_1(t)$. With this in mind, the formula developed in Chapter 2 can be rearranged:

$$\widetilde{r}(t) = \widetilde{x_2}(t) * h_{21}(t) + \widetilde{x}_1(t) * h(t) + n(t). \tag{3.2}$$

The adaptive filter needs to mimic $h(t)$ by processing $\widetilde{x}_1(t)$. If we call the response of the adapted filter $\widehat{h}(t)$, the received signal after the self interference signals have been removed is given by

$$\widetilde{x}_2'(t) = \widetilde{x_2}(t) * h_{21}(t) + \widetilde{x}_1(t) * h(t) - \widetilde{x}_1(t) * \widehat{h}(t) + n(t). \tag{3.3}$$

If $\widehat{h}(t)$ exactly equals to $h(t)$ then the received signal becomes

$$\widetilde{x}_2'(t) = \widetilde{x_2}(t) * h_{21}(t) + n(t). \tag{3.4}$$

In this case an adaptive equalizer can be used to recover $\widetilde{x}_2(t)$ from the channel response and the transmitted data can be recovered [18]. A block diagram illustrating the echo canceller removing the interference on the received signal is shown in Figure 1.7 on page 10.

The adaptive algorithms are never perfect, and there is likely to be come residual interference left over after the cancellation. Using the distributive property of the convolution we can develop an equation for this residual leakage

$$\begin{aligned} \widetilde{x}_2'(t) &= \widetilde{x_2}(t) * h_{21}(t) + \widetilde{x}_1(t) * (h(t) - \widehat{h}(t)) + n(t) \\ &= \widetilde{x_2}(t) * h_{21}(t) + \widetilde{x}_1(t) * h_{residual}(t) + n(t) \end{aligned} \tag{3.5}$$

If for some reason $h(t)$ changes suddenly then $h_{residual}(t)$ will increase in magnitude and the error will increase. Overtime $\widehat{h}(t)$ will adapt to $h(t)$ and the error will return to its steady state level until the channel response changes again.

3.2 Channel Coherence Time

The coherence time of a channel is the length of time over which the channels characteristics are considered to be constant. Coherence time of the channel has been widely studied in wireless communication applications, and those concepts can be extended to wired channels.

There are many ways to define the coherence time of a channel in wireless applications, but the method used for this thesis differs slightly from these. To define the coherence time of a channel consider a channel impulse response at time t_1, $h(t, t_1)$. A transmitted signal, $x(t)$, received at time t_1 would be received as the convolution $y_1(t - t_1) = x(t - t_1) * h(t, t_1)$. Now consider the same channels impulse response but at time t_2. The received signal at this time would be the convolution $y_2(t - t_2) = x(t - t_2) * h(t, t_2)$. For a completely unchanging channel the received signal $y_2(t)$ would be a shifted version of $y_1(t)$ and the system would be considered time invariant. In this case the difference $h(t, t_1) - h(t, t_2)$ would be negligible and the coherence time of the channel would be defined as

$$T_c = t_2 - t_1, \tag{3.6}$$

where T_c is the coherence time of the channel in seconds.

The echo canceller will have a dwell time, T_d, associated with it. The dwell time is here on defined as the time that it takes the echo canceller to adapt to the channel. The maximum possible T_d will depend on the conditions of the cable plant, but the exact T_d will depend on the algorithm used by the chip manufacture as well as a number of other factors [17]. If $T_d > T_c$ then the residual self interference signal will significantly degrade the SNR of the received signal. It is possible for the dwell time of the echo canceller to be several seconds long, which means a channel with a "long" coherence time can still create issues.

This issue of coherence time versus echo canceller dwell time has real effects on the performance of a communication system. It has been shown for wireless communications that the BER (Bit Error Rate) for a RLS (Recursive Least Squares) echo canceller increased

as the coherence time of the channel decreased [19]. The coherence time in a wired channel in the past has generally been assumed to be large enough that this should not be issue. However, in a full duplex network the echo canceller is required to lower the self interference signal by up to 70 dB to achieve the required SNR. If $10\log(\frac{\int h^2(t)dt}{\int(\widehat{h}(t)-h(t))^2dt})$ goes below this 70 dB, then the echo canceller was not able to keep up to the changing channel and the noise floor of the upstream signal will raise.

There are a number of factors that may cause a sudden change in the channel response. Many of the coaxial cables used to transmit to the subscriber households are exposed to the environment, and can be heated by the sun, or shaken by the wind causing the velocity of propagation through the cable to change creating a time variant channel response. This could change the channel response at a rate that the echo canceller could not track. Another cause of channel response change is sudden shifts in the load on the network causing the network reflection coefficients to change,= leading to echoes that can cause a rapidly shifting channel response. All of these scenarios lead to stress on the performance of the echo canceller and unreliable upstream reception.

3.3 Phase Noise and its Limitation on Echo Suppression

In the previous sections it has been argued that the performance of the echo canceller in a full duplex node is susceptible to a rapidly changing channel. In this section it is shown that the phase response of the coaxial cable network is its most unstable characteristic and has the largest effect on the noise floor of the received signal.

Phase noise is generally considered a characteristic inherent to oscillators in which random fluctuations in the phase of the generated signal lead to spectral spreading. Phase noise on a sinusoid can be represented in the following form

$$x(t) = \sin(2\pi ft + \boldsymbol{\phi}(t)), \tag{3.7}$$

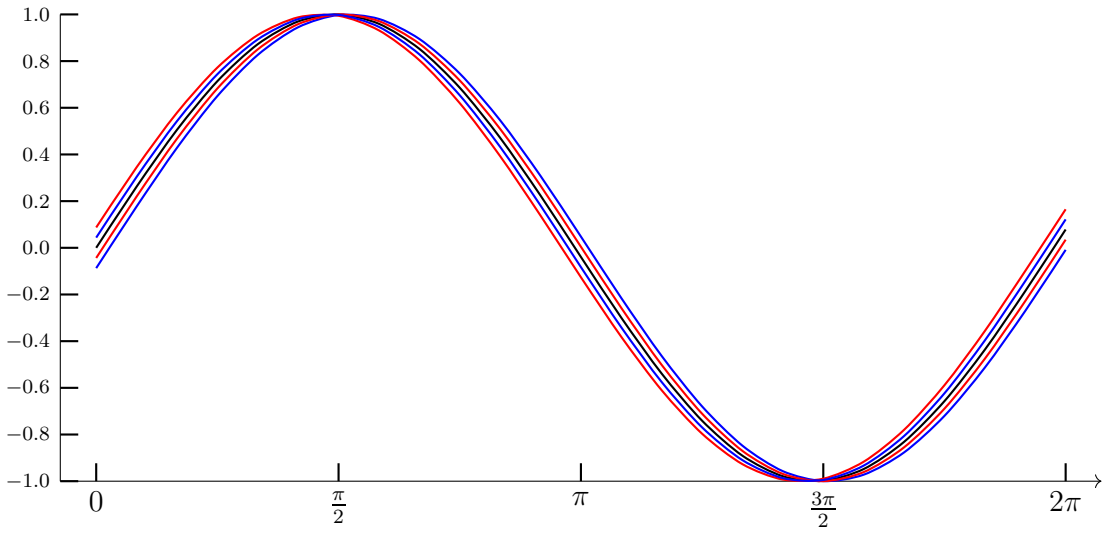

Figure 3.3 Sinusoidal Signal with Phase Noise

where $\boldsymbol{\phi}(t)$ has units radians and represents random fluctuations in the phase of the signal as a function of time. An example of a sinusoid with phase noise is illustrated in Figure 3.3, in this figure several cycles of a sinusoid corrupted by phase noise are shown on the same graph to illustrate the fluctuations in the phase of the sinusoid over time.

To determine the effect that random phase fluctuations has on the signal, we can use the sum difference trig identities to rewrite the sinusoid with phase noise as

$$\begin{aligned} x(t) &= \cos(2\pi ft + \boldsymbol{\phi}(t)) \\ &= \cos(2\pi ft)\cos(\boldsymbol{\phi}(t)) - \sin(2\pi ft)\sin(\boldsymbol{\phi}(t)). \end{aligned} \tag{3.8}$$

Using the first order Maclauren approximations [20] for $\cos(x) \approx 1 - \frac{x^2}{2} \approx 1$ shows for small x that $cos(x) \approx 1$ and $\sin(x) \approx x$. Assuming $|\boldsymbol{\phi}(t)| < 0.1$ radians, (3.8) simplifies to

$$x(t) = \cos(2\pi ft) - \boldsymbol{\phi}(t)\sin(2\pi ft). \tag{3.9}$$

Equation 3.9 shows that the phase noise can be modelled as additive noise consisting of a quadrature sinusoid of the same frequency modulated by the phase noise itself. In the

frequency domain phase noise appears as a kind of phase modulation and as such causes spreading in the impulse that represents a pure sinusoid in the frequency domain. An example of this spectral spreading is shown in Figure 3.4. The delta function at $F = F_c$ represents a pure sinusoid while the sagging lines represent a sinusoid with phase noise.

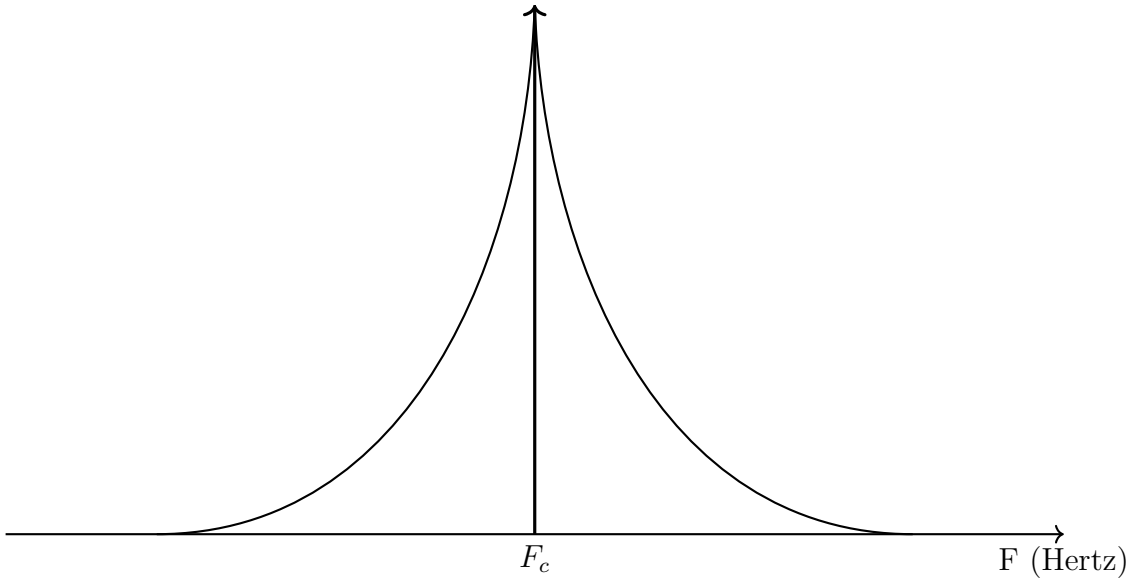

Figure 3.4 Spreading of a Tone's Spectrum Caused by Phase Noise

Consider the simplest case of a single tone being transmitted through a channel and the echo canceller attempting to remove it to determine the effect that phase fluctuations have on the echo canceller. For this case it is assumed that any amplitude noise is negligible and that the echo canceller has perfectly adapted to the channel before the phase fluctuations begin. In this scenario after the self interference subtraction the output is given by

$$e(t) = \cos(2\pi ft + \boldsymbol{\phi}(t)) - \cos(2\pi ft) \tag{3.10}$$

Using the same approximations used in (3.9) the error term $e(t)$ can be rewritten as:

$$e(t) \approx -(\cos(2\pi ft)\frac{\boldsymbol{\phi}^2(t)}{2} + \boldsymbol{\phi}(t)\sin(2\pi ft)) \tag{3.11}$$

The power in this error term can then be calculated by taking the squared mean of the result.

$$\overline{e^2(t)} \approx \frac{\overline{\boldsymbol{\phi}^4(t)}}{8} + \frac{\overline{\boldsymbol{\phi}^2(t)}}{2} \tag{3.12}$$

Since $\boldsymbol{\phi}(t)$ will be very small and $\boldsymbol{\phi}^4(t) << \boldsymbol{\phi}^2(t)$, the high order terms can be ignored and the above equation simplified to $\overline{e^2(t)} \approx \frac{\overline{\boldsymbol{\phi}^2(t)}}{2}$.

With the error power found it is possible to calculate the SNR of the received signal. In this case the SOI (Signal of Interest) is a single sinusoid with unit amplitude, the power in such a signal is $0.5V^2$. The SNR of a single sinusoid at the output of the echo canceller can be calculated as follows

$$\begin{aligned} \text{SNR} &= 10\log_{10}\left(\frac{1/2}{\frac{\overline{\boldsymbol{\phi}^2(t)}}{2}}\right) \\ &= 10\log_{10}\left(\frac{1}{\overline{\boldsymbol{\phi}^2(t)}}\right). \end{aligned} \tag{3.13}$$

The result of (3.13) clearly shows how the SNR of our received signal will decrease as the magnitude squared of the phase fluctuations increases.

It is of worth to discuss what exactly is meant by phase fluctuations at this point. The phase fluctuation given by $\boldsymbol{\phi}(t)$ refers to how much the phase response of the channel varies from its original state at a given time t. A perfectly static channel would have $\boldsymbol{\phi}(t)$ equal to 0 as its phase response stays the same at all times. A varying channel would have $\boldsymbol{\phi}(t)$ as non zero as at some time t the phase response of the channel is not the same as it was at $t = 0$. So when it is discussed that the phase of the channel is varying it is meant that the phase response taken at time $t = 0$ is not the same as when it is taken at time $t = t_1$ and the parameter $\boldsymbol{\phi}(t)$ describes how much it has changed in that time frame in cycles.

The analysis until this point has ignored the effect of amplitude noise on the echo cancellation. The amplitude noise has been ignored because the deviations in the amplitude

relative to the deviations in phase are less impactful and some what mitigable when compared to the effects of the phase noise [21]. We repeat the above analysis with the change of adding the term $\boldsymbol{a}(t)$ to the amplitude to see the effect of amplitude noise in tandem with phase noise.

$$e(t) = (1 + \boldsymbol{a}(t)) \cos(2\pi f t + \boldsymbol{\phi}(t)) - \cos(2\pi f t); \tag{3.14}$$

Repeating the analysis done in (3.12) we can write an expression for the power in the error signal taking the amplitude noise into account.

$$\overline{e^2(t)} = \frac{\overline{\boldsymbol{a}^2(t)}}{2} + \frac{\overline{\boldsymbol{a}^2(t)\boldsymbol{\phi}^4(t)}}{8} + \frac{\overline{\boldsymbol{\phi}^2(t)\boldsymbol{a}^2(t)}}{2} + \frac{\overline{\boldsymbol{\phi}^4(t)}}{8} + \frac{\overline{\boldsymbol{\phi}^2(t)}}{2} \tag{3.15}$$

The cross terms can be completely ignored as they will be insignificantly small. Removing the cross terms leaves the same equation found in (3.12) with the addition of $\frac{\overline{\boldsymbol{a}^2(t)}}{2}$. As the amplitude noise is relatively stable when compared to the phase noise $\overline{e^2(t)}$ will be dominated by the phase terms and the amplitude terms can be ignored.

The fluctuations in the phase response of the network have a direct effect on the SNR of the received signal in the full duplex node. The phase noise in the clock circuitry used to drive the DAC and the ADC also lower the SNR as explained in Chapter 2. In full duplex wireless communications it has been demonstrated that the phase noise in the oscillator used to up convert and down convert was the limiting factor on the self interference cancellation [10]. To properly characterize the full duplex network it is therefore important to determine the stability of the phase response of the network in the full duplex band.

3.4 Phase Measurement Requirements

In the previous sections of this chapter it has been shown that there is a need to measure the fluctuations in the phase response of the cable plant in a full duplex DOCSIS network. Furthermore, in (3.13) it was shown that very small fluctuations in the phase response

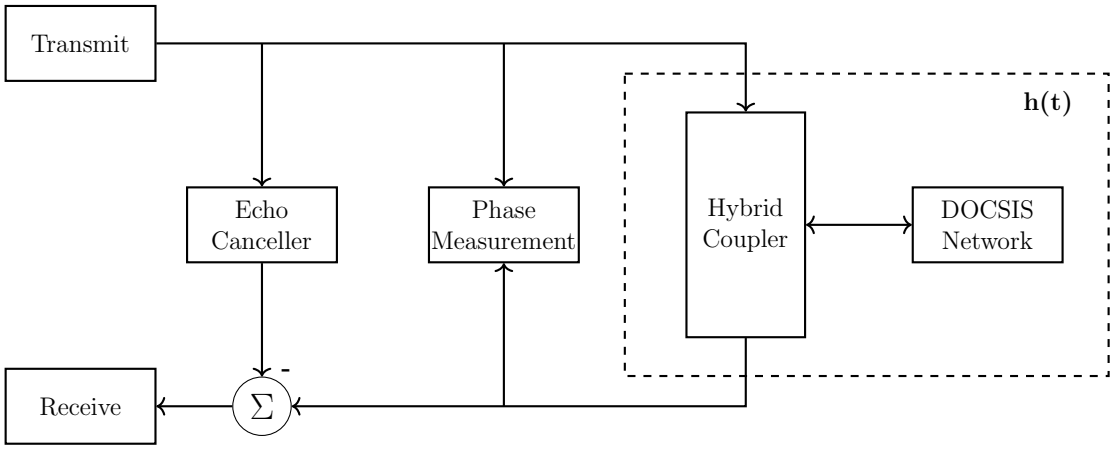

Figure 3.5 Phase Stability Measurement

of the network can lead to temporarily large errors as the echo canceller re-adapts. The requirements of the echo canceller and (3.13) are now used to determine the requirements on a phase fluctuation measurement device.

The diagram shown in Figure 3.5 shows a typical echo canceller with the addition of a new block that compares the phases of the transmit and receive sides of the node. The requirements of this block are determined by assuming that the power in the transmitted signal relative to the variance of the phase variation needs to be 70 dB. That is, it is assumed that the SNR for a sinusoid with an amplitude of 1 must be 70 dB where the SNR is given by

$$\begin{aligned} SNR &\approx 10\log_{10}\left(\frac{1}{\overline{\phi^2(t)}}\right) \\ 70 &\approx 10\log_{10}\left(\frac{1}{\overline{\phi^2(t)}}\right) \\ 10^{70/10} &\approx \frac{1}{\overline{\phi^2(t)}} \\ 10^{-7} &\approx \overline{\phi^2(t)}\ \text{rad}^2 \\ \sqrt{\overline{\phi^2(t)}} &\approx 3.16 \cdot 10^{-4}\ \text{rad}. \end{aligned} \tag{3.16}$$

From the results in (3.16), it can be seen that the phase measurement system must be able to

measure deviations in phase of approximately 300 μrad. Ideally the system would be able to measure even smaller deviations than 300 μrad to give a better understanding of the effects of different disturbance events on the network. A graph outlining the relation between the SNR and the RMS of the phase variance is shown in Figure 3.6.

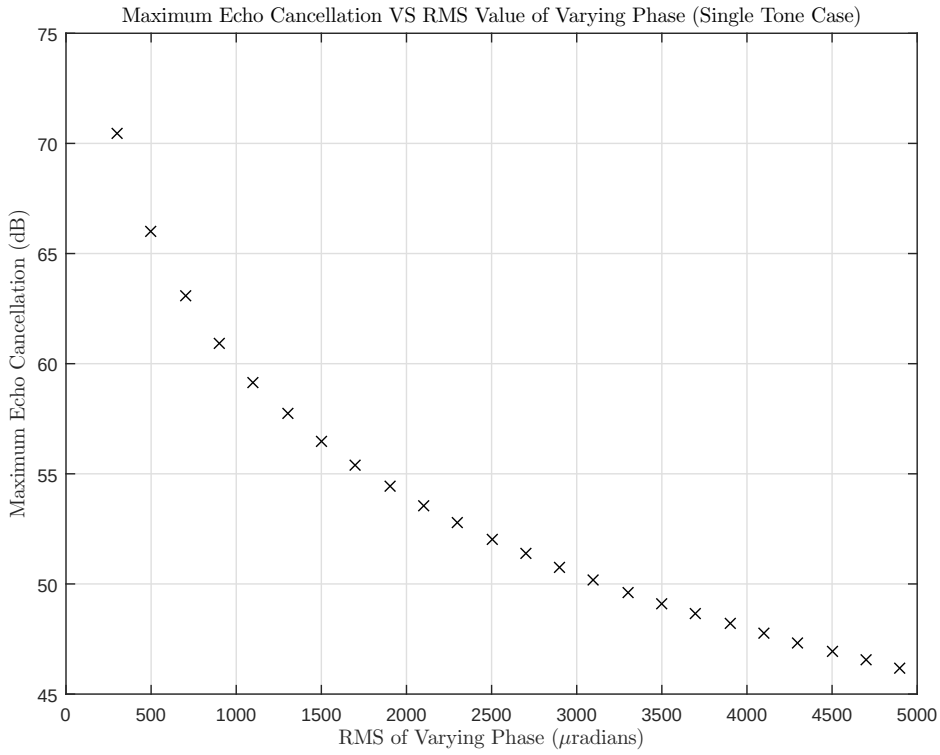

Figure 3.6 RMS of Phase Fluctuations Effect on SNR

As well as being able to measure extremely small deviations in the phase response of the network, the phase measurement device must also be able to store these deviations for later analysis. The device needs to be able to store the measurements over the duration of several T_d of the echo canceller. The device also must be able to store this time domain phase deviation information with reasonable resolution, i.e. the data must be sampled at a high enough rate to capture the phase deviations.

Upon initial examination it may seem like a VNA (Vector Network Analyzer) is suitable for making these measurements, but VNAs are unsuitable for a number of reasons. A VNA

is a powerful device capable of measuring among other things the S parameters of a network. Since the S parameters are not immediately applicable to the measurements being made in this thesis, the VNA is not the right tool. VNA also have a limited aperture time, which is the time interval over which the measurement is made. The VNA has to convert its frequency domain measurements to time domain which means that it is very possible some of the micro fluctuations in phase would be lost in conversion. Finally, since the measurement device needs to be able to measure fluctuations on the order of μradians a very high end VNA would have to be used and even then it is unlikely to be appropriate due to the other limitations mentioned [22] [23].

Since the VNA is unsuitable for making these measurements, a method must be developed that meets all of the requirements outlined in this section. The development of such a measurement scheme is the subject of Chapter 4.

4. Phase Fluctuation Measurement With a Phase Locked Loop

The subject of this chapter is to explore a Phase Locked Loop (PLL) based method for measuring the phase fluctuations of a full duplex DOCSIS channel. While other methods exist [24] [25] the PLL is chosen for reasons explained in this chapter. A theoretical basis for the approximations and error inherent to the system is provided and a specific type of PLL is selected. This chapter focuses on theoretical analysis, the implementation of the PLL is left to Chapter 5.

4.1 Phase Locked Looped Quadrature Method

A phase locked loop is a circuit that is commonly used to lock the phase of one sinusoid to another. It is also used to synthesize clock signals, synchronize data signals, stabilize oscillators, reduce timing jitter, and even in the control of the angular velocity of electric motors [26]. A phase locked loop is a feedback control system, consisting of gains, accumulators, and loop filters. A generic digital phase locked loop is shown in Figure 4.1.

The generic PLL can be a Type 1, Type 2, or Type 3 loop depending on the constants K_1, K_2, and K_3. A Type 1 PLL is able to track the frequency of the incoming signal with zero error, a Type 2 PLL is able to track the frequency and phase of the incoming signal with zero error, and a Type 3 PLL is able to track a linearly changing frequency and the phase of an incoming signal [26]. For reasons explained in this chapter a Type 2 PLL is selected for the measurements to be made.

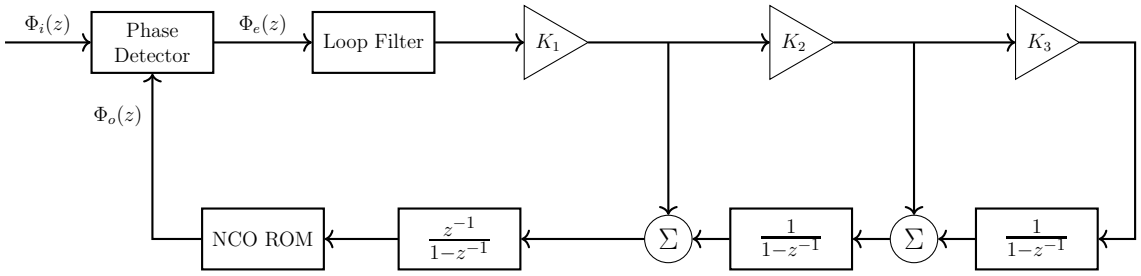

Figure 4.1 Generic Digital Phase Locked Loop Circuit

The function of the PLL is to produce a noise free sinusoid of the same frequency and phase as the noisy sinusoidal input. The PLL controls the frequency and phase of the output sinusoid by driving the difference in phase between the two signals, $\Phi_e(z)$, to zero. It does this by comparing the phase of the generated sinusoid to the phase of the reference sinusoid with a phase detector as illustrated in Figure 4.1. Eventually, the two sinusoids are locked in quadrature with each other, which means the sinusoid at the output of the NCO has exactly the same frequency as the reference, but differs in phase by $\pi/2$. The phase locked loop shown in Figure 4.1 is a Type 3 circuit if all three gain parameters K_1, K_2, K_3 are non zero. This Type 3 circuit is converted to a Type 2 by setting K_3 to 0, and it is converted to a Type 1 by setting both K_2 and K_3 to 0.

The crux of the phase locked loop is the phase detector. It produces an output which is proportional to the difference between the phase of the reference and NCO output sinusoids. While the function of a phase detector is well defined, there is no circuit that can implement a perfect phase detector. A multiplier can be used, but it generates a double frequency term that must be treated as noise. The output of the multiplier phase detector is modeled as a difference in phase with additive double frequency error.

It is understood in Figure 4.2 that a ϕ denotes a phase signal in cycles and a ϕ^r denotes a phase signal in radians. For the Type 1 PLL as shown in Figure 4.2 the output of the multiplier phase detector is given by

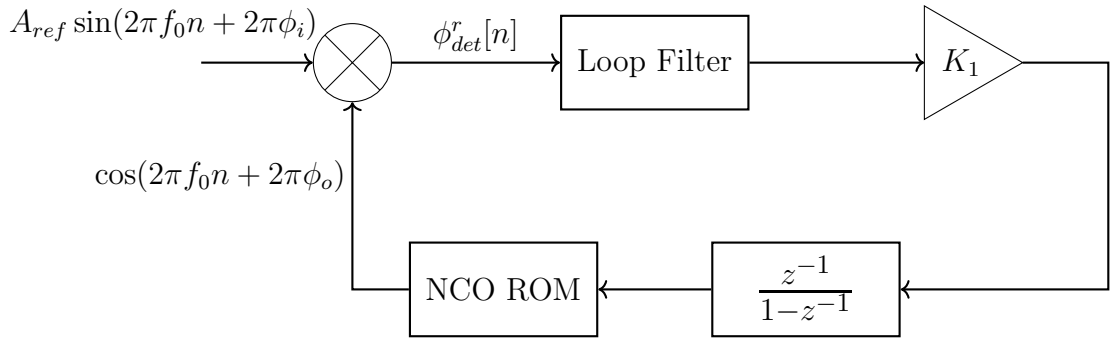

Figure 4.2 Type 1 Phase Locked Loop with Multiplier Phase Detector

$$\begin{aligned} 2\pi\phi_{det}[n] &= A_{ref}\sin(2\pi f_0 n + 2\pi\phi_i[n]) \cdot \cos(2\pi f_0 n + 2\pi\phi_o[n]) \\ &= \frac{A_{ref}}{2}\sin(2\pi(\phi_i[n] - \phi_o[n])) + \frac{A_{ref}}{2}\sin(4\pi f_0 n + 2\pi(\phi_o[n] + \phi_i[n])) \\ &\approx \frac{A_{ref}}{2}(2\pi(\phi_i[n] - \phi_o[n])) + \frac{A_{ref}}{2}\sin(4\pi f_0 n + 2\pi(\phi_o[n] + \phi_i[n])). \end{aligned} \tag{4.1}$$

The double frequency component arising from the multiplication can be suppressed with a filter, which is the purpose of the "loop" filter in Figure 4.2. In this case the detector output is $\frac{A_{ref}}{2}$ times the difference in phase between the NCO output and the reference input. Therefore the multiplier detector is said to have a gain, K_d, of $\frac{A_{ref}}{2}$.

The quality of the approximation $K_d\sin(2\pi(\phi_i[n]-\phi_o[n])) \approx K_d(2\pi(\phi_i[n]-\phi_o[n]))$ is now explored. It is based on a first order Maclaurin approximation of the sine function shown in (4.2).

$$\begin{aligned} \sin(x) &= \sum_{k=0}^{\infty} \frac{(-1)^k}{(2k+1)!}x^{2k+1} \\ &= x - \frac{x^3}{3!} + \frac{x^5}{5!} - ... \end{aligned} \tag{4.2}$$

The above series expansion shows that for $x << 1$ the function $\sin(x)$ is well approximated by x, due to the the higher order terms becoming increasingly insignificantly small. Of interest is the point at which the error in the approximation becomes unacceptable. If the highest acceptable phase difference in units radians is called ϕ^r_{max} it is of interest to determine at

what point the ratio of the approximate to the true value falls below some acceptable value of $\mathscr{E}_{approx}$.

$$\mathscr{E}_{approx} = \frac{\sin(\phi^r_{max})}{\phi^r_{max}} \tag{4.3}$$

Assuming the 3rd order Maclaurin approximation is much more accurate than the 1st order approximation it can be used for $\sin(\phi_{max})$ in (4.3) and the equation becomes

$$\mathscr{E}_{approx} = \frac{\phi^r_{max} - \frac{(\phi^r_{max})^3}{6}}{\phi^r_{max}}. \tag{4.4}$$

Rearranging (4.4) yields

$$\phi^r_{max} = \sqrt{6(1 - \mathscr{E}_{approx})}. \tag{4.5}$$

From (4.4) an acceptable error ratio of 0.95 restricts the largest phase difference to less than 0.54771 radians. This constraint exists on top of the absolute constraint of $\phi^r_{max} < \frac{\pi}{2}$, as the detector can not detect a phase difference outside the interval $(-\frac{\pi}{2}, \frac{\pi}{2})$ radians.

The output of the ideal phase detector in the phase locked loop circuit of Figure 4.1 is the difference in phase between the reference sinusoid and the locally generated output sinusoid. If the reference is a sinusoid of fixed and known frequency the loop in Figure 4.2 eventually drives this difference to zero and the NCO output becomes locked. The goal of this circuit is to measure the phase variations in the channel response of a DOCSIS network.

The phase locked loop circuits shown so far can measure phase variations on a single sinusoid being received from some network. If the response time of the phase locked loop is designed to be less than that of the self interference canceler, then phase changes too sudden for the self interference canceler to immediately deal with will certainly appear on the output of the phase detector in a phase locked loop that is locked. This means that examination of

the output of the phase detector in a locked phase locked loop will track the variations in phase of the channel response. An example of a measurement system is shown in Figure 4.3.

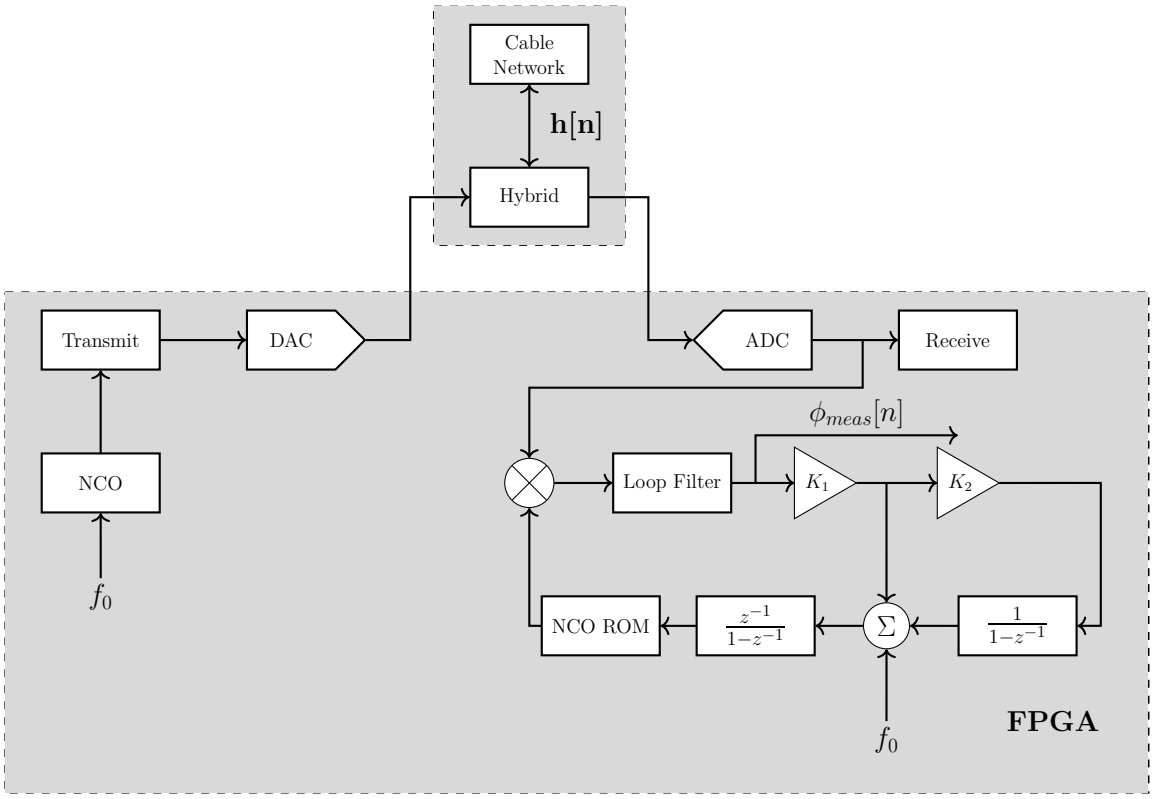

Figure 4.3 Hardware Setup for a Phase Locked Loop Based Measurement System

The Phase Locked Loop circuit shown in Figure 4.3 is a Type 2 circuit modified to have a second input for frequency. The reason for this frequency input will be explained in detail later on in this chapter, but for now it is enough to say that this greatly decreases the time it takes for the PLL to lock to the incoming sinusoid. The Phase Locked Loop in Figure 4.3 is a Type 2 for reasons that will also be stated later on in this chapter, but for now it is enough to state that a Type 2 circuit is able to track both changes in phase and frequency of the input sinusoid. The output of this circuit shown as $\phi_{meas}[n]$ is the fluctuations in the phase response in cycles of the channel as a function of time, but scaled by K_d. If desired, the PLLs feedback circuit can be temporarily disconnected so that the loop stops adapting to changing phase. The circuit will no longer attempt to force the phase difference to zero, and will directly output the phase with respect to the reference from the time the feedback

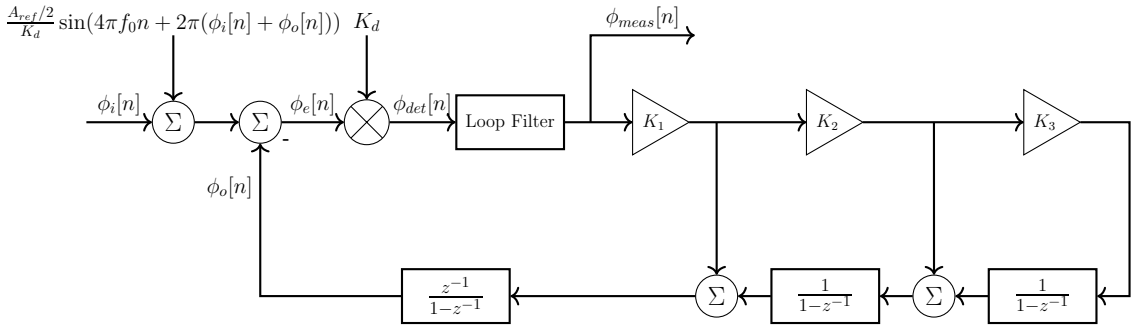

Figure 4.4 Equivalent Model of a Type 3 Phase Locked Loop

was disconnected.

4.2 Output of a Phase Locked Loop

In the previous section the phase detector of the PLL circuit was a multiplier and its operation was analyzed using a small signal approximation. In this section we will take a step back and model the phase detector as a simple subtraction followed by the detector gain of K_d with the double frequency term arising from the multiplication shown as additive noise injected before the subtraction. For this analysis, the input to the system $\phi_i[n]$, i.e. the reference input, is the phase of the reference signal in units cycles. The output phase, $\phi_o[n]$, is also a phase signal of unit cycles. The direct output of the subtraction of the reference and output phase, $\phi_e[n]$ is the error, which is of units cycles. This equivalent system is shown in Figure 4.4 for a generic Type 3 PLL. It is again noted that this PLL can be converted to a Type 1 or a Type 2 circuit by setting the appropriate K values to zero.

The model shown in Figure 4.4 is a phase model of Figure 4.1. It differs from the circuit in Figure 4.1 in a few ways. The NCO ROM is not part of the phase model as it now deals with only phase signals. The additive noise from the double frequency component at the output of the phase detector multiplier is modelled as additive phase noise injected on the input side the phase detector. The additive phase detector noise and the output of the phase detector is scaled by K_d so that the output of this circuit is equivalent to the output of Figure 4.1. The following equations prove the equivalence of these circuits

$$
\begin{aligned}
2\pi\phi_{meas}[n] &= (2\pi\phi_i[n] + \sin(4\pi f_0 n + 2\pi(\phi_i[n] + \phi_o[n])) - 2\pi\phi_o[n]) \cdot K_d * h_{lf}[n] \\
&= (K_d 2\pi(\phi_i[n] - \phi_o[n]) + K_d \sin(4\pi f_0 n + 2\pi(\phi_i[n] + \phi_o[n]))) * h_{lf}[n] \\
&= (K_d(2\pi\phi_i[n] - 2\pi\phi_o[n]) + K_d \sin(4\pi f_0 n + 2\pi(\phi_i[n] + \phi_o[n]))) * h_{lf}[n] \\
&\approx K_d(2\pi(\phi_i[n] - \phi_o[n])),
\end{aligned}
\tag{4.6}
$$

From (4.6) it is clear that the output of this equivalent circuit matches the output of the PLL found in the previous section. The advantage of this equivalent circuit is that the transfer functions to any point in the circuit can now be developed with much greater ease.

4.3 Transfer Function Analysis of a Phase Locked Loop

It is possible to generate a frequency response for any type of PLL from the circuit shown in Figure 4.4 using the approximation given in (4.6). With resource utilization in mind, the analysis begins with the smallest possible circuit, which is a Type 1 PLL. Should the analysis show a Type 1 PLL is inadequate, the type of the PLL will be incremented until a circuit is found that is able to make the measurements with acceptable error.

4.3.1 Analysis of a Type 1 PLL

To generate the frequency response of a Type 1 PLL the circuit in Figure 4.4 is again simplified. The circuit now neglects the double frequency additive noise component and shows only the reference phase $\phi_i[n]$ in cycles. The loop filter shown in previous versions of this circuit is now represented by a generic transfer function, $H_{lf}(z)$, to simplify analysis.

From this circuit it is now easy to find the transfer functions that characterize a Type 1 PLL. A circuit diagram of a Type 1 PLL is shown in Figure 4.5. To begin developing the closed loop and error transfer functions the open loop transfer function, $G_1(z)$, is needed and is defined as follows

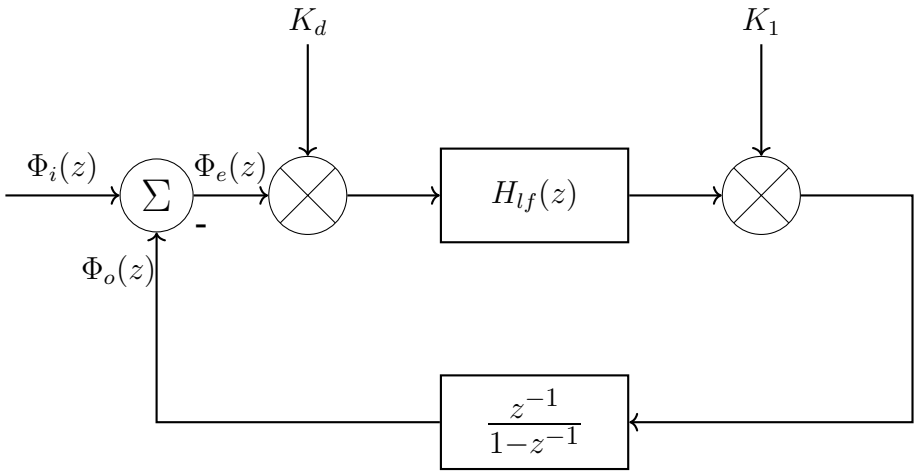

Figure 4.5 Type 1 Phase Locked Loop

$$\begin{aligned} G_1(z) &= \frac{\Phi_o(z)}{\Phi_e(z)} \\ &= \frac{K_d \cdot H_{lf}(z) \cdot K_1 \cdot z^{-1}}{1 - z^{-1}}. \end{aligned} \tag{4.7}$$

With the open loop transfer function defined, the closed loop and error transfer functions are easily found. The closed loop transfer function, $H_1(z)$ is shown in (4.8). To help the readability of the equations the substitution $K_d \cdot H_{lf}(z) \cdot K_1 = R_1(z)$ has been made,

$$\begin{aligned} H_1(z) &= \frac{\Phi_o(z)}{\Phi_i(z)} \\ &= \frac{G_1(z)}{1 + G_1(z)} \\ &= \frac{R_1(z) \cdot z^{-1}}{1 - (1 - R_1(z)) \cdot z^{-1}}. \end{aligned} \tag{4.8}$$

Lastly, the equation for the error transfer function is found using the formulas for the open and closed loop transfer functions,

$$
\begin{aligned}
E_1(z) &= \frac{\Phi_e(z)}{\Phi_i(z)} \\
&= 1 - H_1(z) \\
&= \frac{1}{1 + G_1(z)} \\
&= \frac{1 - z^{-1}}{1 - (1 - R_1(z)) \cdot z^{-1}}.
\end{aligned} \tag{4.9}
$$

There are some immediate observations that can be made about the transfer functions found in (4.7), (4.8), and (4.9). The first observation that can be made is that the error transfer function, $E_1(z)$, has a high pass response with a zero at $\omega = 0$. Another observation is that the closed loop transfer function, $H_1(z)$, has a low pass response with a pole at $z = 1 - K_1$. For the purposes of the measurements to be made with this circuit the transfer function of the most interest is $E_1(z)$ as its high pass type response filters out the low frequency changes in the phase that the echo canceller can deal with and passes the high frequency phase changes that are of interest. The closed loop transfer function does the opposite, $H_1(z)$ would only pass the low frequency changes that the echo canceler can properly track anyway and filter out the fluctuations of interest.

The incoming phase in cycles, $\phi_i[n]$, for a single tone input would be a ramp function, meaning that the PLL must be able to track a ramp input with no error. To find the steady state error for a ramp input in a Type 1 PLL we use a combination of the error transfer function, $E_1(z)$, and the final value theorem which is given for the general case by

$$
\lim_{n \to \infty} f[n] = \lim_{z \to 1} (1 - z^{-1}) \cdot F(z), \tag{4.10}
$$

setting $\Phi_i(z) = \frac{\alpha \cdot z^{-1}}{(1 - z^{-1})^2}$ as the z-transform of a ramp signal with slope α the steady state error in the Type 1 PLL is given by:

$$\begin{aligned}\phi_e[\infty] &= \lim_{z\to 1}(1-z^{-1})\cdot E_1(z)\cdot \Phi_i(z) \\ &= \lim_{z\to 1}(1-z^{-1})\cdot \frac{1-z^{-1}}{1-(1-R_1(z))\cdot z^{-1}}\cdot \frac{\alpha\cdot z^{-1}}{(1-z^{-1})^2} \\ &= \frac{\alpha}{R_1(z)}\end{aligned} \tag{4.11}$$

Since $\phi[\infty] = \frac{\alpha}{R_1(z)}$, the Type 1 PLL shown in Figure 4.5 would not be able to track a ramp input with zero error. The result of the final value theorem shows that the Type 1 PLL would track the slope of the incoming ramp signal correctly, but there would be a constant offset on the error output. This "constant" error is unacceptable for the measurements that need to be made. For this reason a Type 1 PLL as shown cannot be used and we must explore a Type 2 PLL.

4.3.2 Analysis of a Type 2 PLL

As shown in the previous subsection a Type 1 PLL is inadequate. The next logical step is to explore a Type 2 PLL. A Type 2 PLL with the same modifications made to the Type 1 PLL is shown in Figure 4.6.

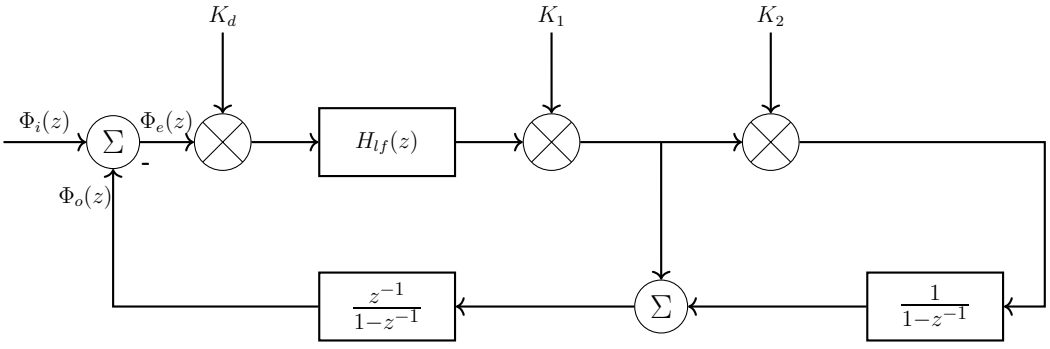

Figure 4.6 Type 2 Phase Locked Loop

This circuit adds a second gain parameter, K_2, and a second accumulator in the feedback path. The extra components in this circuit has the effect of raising the order of the transfer function which is why this circuit is referred to as a Type 2 PLL. The raising of the order of the transfer function of the circuit means that the circuit can track higher order inputs

without error. The open loop transfer function of the Type 2 PLL, $G_2(z)$, is as follows

$$\begin{aligned} G_2(z) &= \frac{\Phi_o(z)}{\Phi_e(z)} \\ &= \frac{R_1(z) \cdot z^{-1} \cdot (K_2 + 1 - z^{-1})}{(1 - z^{-1})^2}, \end{aligned} \tag{4.12}$$

where $R_1(z)$ is the same parameter used in the analysis of the Type 1 PLL. The closed loop transfer function is again easily found from the open loop transfer function

$$\begin{aligned} H_2(z) &= \frac{\Phi_o(z)}{\Phi_i(z)} \\ &= \frac{G_2(z)}{1 + G_2(z)} \\ &= \frac{R_1(z) \cdot z^{-1} \cdot (K_2 + 1 - z^{-1})}{(1 - z^{-1})^2 + R_1(z) \cdot z^{-1} \cdot (K_2 + 1 - z^{-1})}. \end{aligned} \tag{4.13}$$

Finally the function of interest, $E_2(z)$, is found below

$$\begin{aligned} E_2(z) &= \frac{\Phi_e(z)}{\Phi_i(z)} \\ &= 1 - H_2(z) \\ &= \frac{1}{1 + G_2(z)} \\ &= \frac{(1 - z^{-1})^2}{(1 - z^{-1})^2 + R_1(z) \cdot z^{-1} \cdot (K_2 + 1 - z^{-1})}. \end{aligned} \tag{4.14}$$

The error transfer function of the Type 2 PLL again has a high pass response with two zeros at $\omega = 0$. The closed loop transfer function again has a low pass type response, meaning it again is not useful for the purposes of the measurements of interest. The final value theorem is again utilized to determine how the systems steady state error responds to a ramp input. The steady state error becomes:

$$\begin{aligned}\widehat{\phi}_e[\infty] &= \lim_{z\to 1}(1-z^{-1})\cdot E_2(z)\cdot \widehat{\Phi}_i(z) \\ &= \lim_{z\to 1}(1-z^{-1})\cdot \frac{(1-z^{-1})^2}{(1-z^{-1})^2 + R_1(z)\cdot z^{-1}\cdot(K_2+1-z^{-1})}\cdot \frac{\alpha\cdot z^{-1}}{(1-z^{-1})^2} \\ &= 0,\end{aligned} \tag{4.15}$$

This means the Type 2 PLL can perfectly track the desired input allowing us to make the channel measurements of interest with no steady state error.

Examination of the Type 1 and Type 2 PLL circuits show that the adder before the $\frac{z^{-1}}{1-z^{-1}}$ accumulator in both circuits is equivalent to the frequency input of an NCO. This means that for the PLL to lock to a particular sinusoid, at a particular frequency and phase, the value after that adder must eventually converge to the value of the frequency of the incoming sinusoid in cycles. This explains the non zero steady state error of the Type 1 PLL, if the value after that adder (before the accumulator) must be the frequency of interest then the error cannot be non zero as it is equal to a scaled version of the error. For the Type 2 PLL the error can be non zero as the second accumulator can converge to the required frequency allowing for the output of the phase detector to be zero. This information is useful because it means that if we know the frequency of the sinusoidal input to the PLL, instead of waiting for the PLL to converge to this frequency we can directly input it. Doing this we can greatly decrease the lock time of the PLL, because it would then only need to adjust the phase of the generated sinusoid. It also means that the Type 1 PLL can be made to have zero steady state error.

If the frequency is input directly into the the Type 1 PLL then the input to the accumulator will be the frequency of interest, meaning that the contribution from the phase detector must be zero. If the contribution from the phase detector is zero then $\phi_e[n]$ is zero and the steady state error is zero. This implies that the Type 1 PLL could be used for the measurements, however this would be under the critical assumption that there is no drift in the frequency of the received sinusoid. If the frequency of the incoming sinusoid and the value directly added to the Type 1 PLL differ at all then the steady state error will again

appear on the output. This can happen due to a number of reasons such as differences in the word format of the frequency input to the transmit NCO and the frequency input to the PLL. For this reason the preferred circuit is still the Type 2 PLL with the frequency input directly as it allows for some error in the frequency input while still decreasing the lock time of the loop.

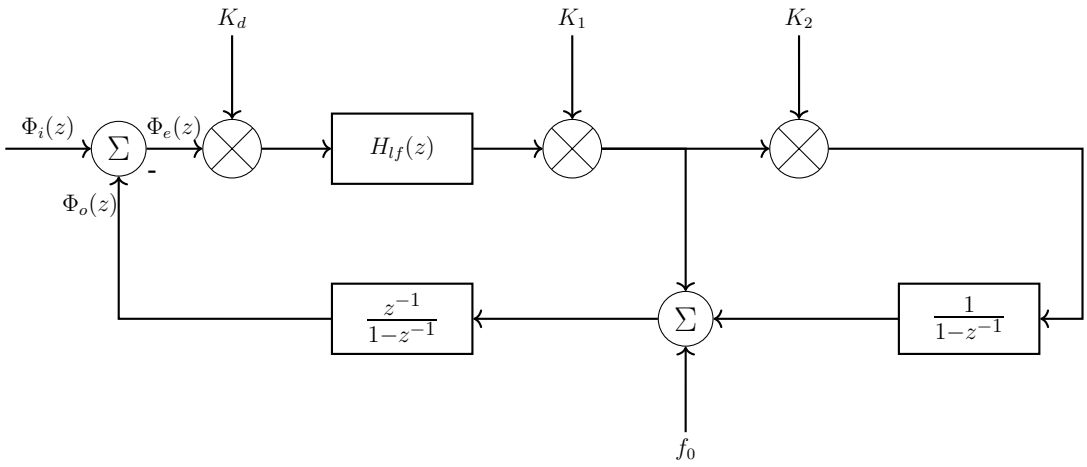

Figure 4.7 Type 2 Phase Locked Loop with Frequency Inputted

5. Implementation of a Phase Locked Loop Circuit on an FPGA

In this chapter the Phase Locked Loop circuit developed in Chapter 4 will be designed to meet the measurement specifications outlined at the end of Chapter 3. The two gain parameters, K_1 and K_2, are found to optimize the transfer function of the PLL. The word formats of the signals in the PLL circuit are carefully chosen to keep the quantization noise within acceptable levels. The objective of the chapter is to generate a detailed schematic diagram of the PLL circuit to be implemented in the FPGA.

5.1 Parameter Selection in a Type 2 Phase Locked Looped

A Type 2 PLL is capable of measuring the type of fluctuations encountered in the phase response of a full duplex channel through proper selection of K_1 and K_2. The Type 2 PLL will be used to determine the bandwidth of the phase fluctuations to determine the rate at which the self interference canceller must adapt. It is not necessary for the PLL to produce the lower frequencies of the phase fluctuations since the echo canceller will certainly be able to cancel them if it can cancel the higher frequencies. A transfer function for the modelled error output, $\Phi_e(z)$, from the input, $\Phi_i(z)$, is found in (4.14) on page 65.

In (4.14) there are a number of parameters that have yet to be defined, namely $R_1(z) = K_d \cdot H_{lf}(z) \cdot K_1$ and K_2. These parameters control the placement of the poles for the error transfer function, which in turn control the shape of the frequency response of the PLL. A_{ref} is the amplitude of the reference sinusoid, which is also referred to as the input sinusoid, and

$K_d = \frac{A_{ref}}{2}$ is the small signal detector gain. The transfer function of the loop filter $H_{lf}(z)$ can be ignored for now as its effect on the shape of $E_2(z)$ should only be to suppress the double frequency component. It is therefore up to the parameters K_1 and K_2 to shape the frequency response.

The error transfer function has a complicated transfer function and it is not simple to choose K_1 and K_2 to get a desired frequency response. The process of selecting these parameters is greatly simplified by constraining $K_2 = K_1/2$. It has been shown that by constraining K_2 in this way the transfer function will have complex poles resulting in a monotone increasing pass band [26]. By constraining the gains in this way, the process of finding them is greatly simplified because it becomes a task of setting the stop band corner frequency and the associated attenuation and solving for K_1.

An example frequency response is shown in Figure 5.1. This example diagram shows the frequency response of the circuit is of high pass type. The assumption was made in Chapter 3 that the echo canceller may have a dwell time, T_D, of up to 10 seconds. This means that any changes in the channel that occur over a period longer than 10 seconds are certain to be tracked by the echo canceller. This also means that any changes with a period less than 10 seconds may not be properly tracked by the echo canceller. For this reason the pass band corner frequency of the error transfer function of the Type 2 PLL is set to $F_{pass} = 1/T_D = 0.1$ Hz. The gain at the pass band corner frequency is specified as -1 dB.

The problem reduces to solving (5.1) to find an appropriate value for K_1. It is noted that in (5.1) it is assumed that $K_d = 1$ and that the effects of $H_{lf}(z)$ are negligible at low frequencies thus simplifying $R_1(z)$ to be equal to K_1.

$$10^{-1/20} = \left| \frac{(1 - e^{-j \cdot \frac{0.1}{F_s} \cdot 2\pi})^2}{(1 - e^{-j \cdot \frac{0.1}{F_s} \cdot 2\pi})^2 + K_1 \cdot e^{-j \cdot \frac{0.1}{F_s} \cdot 2\pi} \cdot (\frac{K_1}{2} + 1 + e^{-j \cdot \frac{0.1}{F_s} \cdot 2\pi})} \right|. \tag{5.1}$$

Setting the sampling frequency $F_s = 48$ MHz and solving this relation numerically using MATLAB results in $K_1 = 1.8179 \cdot 10^{-16}$. It is possible to implement the multiplication by

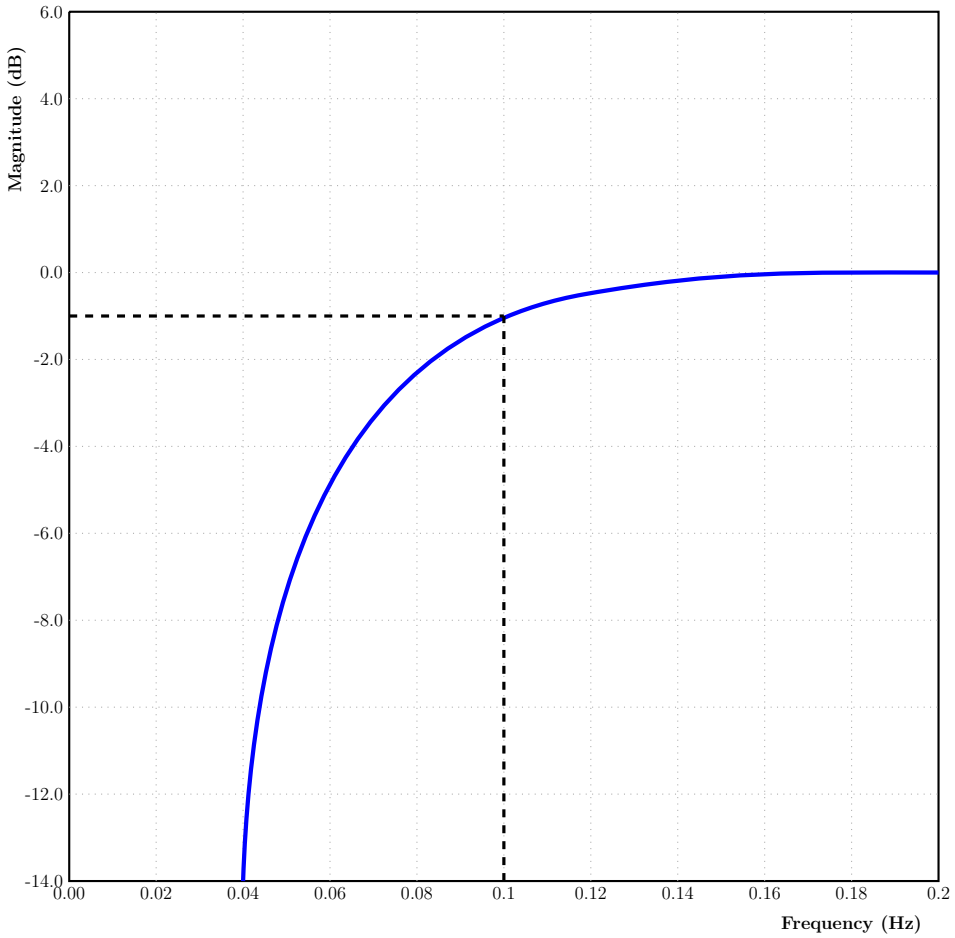

Figure 5.1 Magnitude Response of $E_2(z)$ Showing a Gain of -1 dB at a Stop Frequency of 0.1 Hz

K_1 and K_2 without the use of multiplier hardware if K_1 and K_2 are powers of two. With this constraint the multiplication associated with them can be implemented by bit shifting. Using MATLABs *nextpow2* function it is found that K_1 can be approximated as 2^{-52}, which is equivalent to shifting the bits of the input to the "multiplier" 52 bit positions to the right to get the output.

Determining the convergence time, which is the time from when the PLL is turned on until it reaches steady state, of a phase locked loop circuit is very difficult to predict [26]. However, empirical evidence shows that the convergence time is longer for smaller K_1 and K_2 values [27]. This implies that very small K_1 and K_2 will lead to a longer convergence

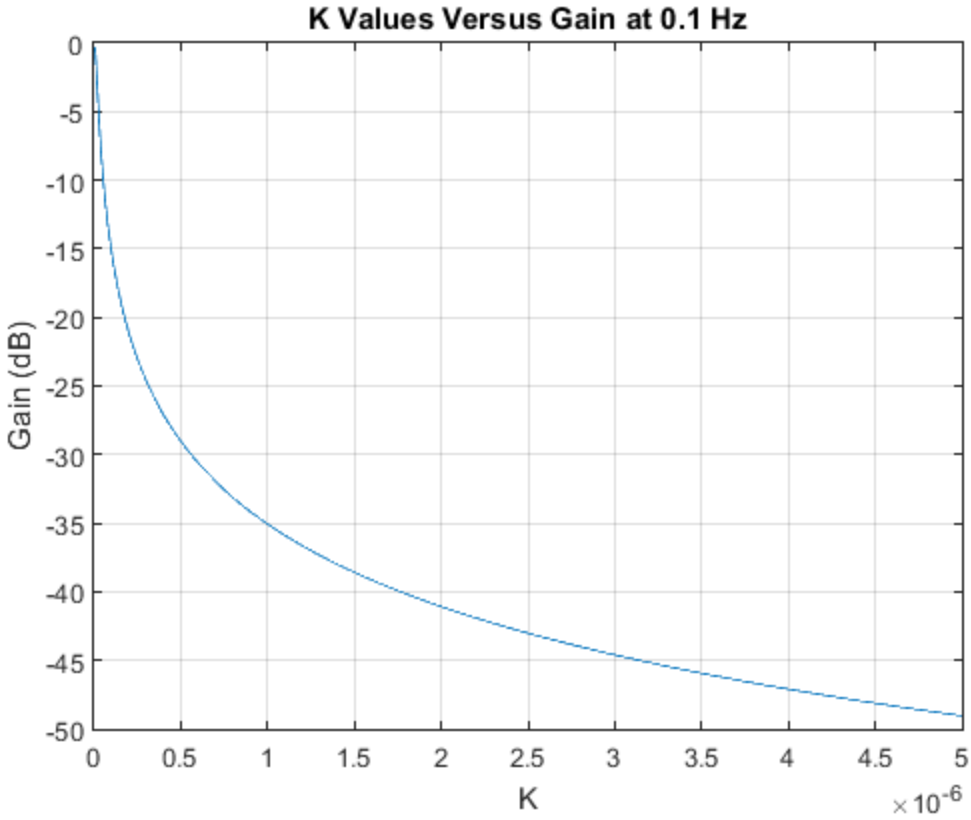

Figure 5.2 Gain at 0.1 Hz for Different Values of K

time than very large K_1 and K_2. Unfortunately, the pass band response dictates that K_1 and K_2 must be small. A circuit has been designed and built to address this issue and is outlined in the next subsection.

5.1.1 Gain Control Circuit

In order to minimize the time it takes for the PLL to converge to the phase of the reference sinusoid it is desirable that the gain parameters, K_1 and K_2, not initialize at the very small values calculated in the previous section. In order to get around this short coming, a circuit has been built that adjusts the gain parameters over time by initializing them at large values and then slowly adjusting them until they converge to their final small values. This process can be thought of as getting the PLL to lock with course values and then once locked slowly move the to the values needed for the narrow bandwidth response. The basic algorithm is

as follows:

- The power in the error output of the phase locked loop is computed by averaging its square over a user defined period, N, in samples $\overline{e^2[n]} = \frac{1}{N}\sum_{n=0}^{N-1} e^2[n]$. The average power is stored, reset, and recalculated every N samples.

- The rate of change of this error power is then calculated by estimating it as a scaled difference, $\frac{\sum_{n=0}^{N-1} e^2[n] - \sum_{n=0}^{N-1} e^2[n-N]}{2N}$. This is accomplished by comparing the current average of N samples to the previous average of N samples at the end of every N sample period.

- If the rate of change of the error is found to be below a user defined threshold, and persists below the threshold for a predetermined number of cycles then the circuit is considered to be in steady state. In this case K_1 and K_2 are instructed to reduce by a factor of two.

- The "rate of change" threshold is divided by two and the above algorithm is repeated until K_1 and K_2 arrive at their final values.

- Once K_1 and K_2 are at their final values the circuit continuously monitors the error output of the PLL. If the rate of change of error exceeds a preset threshold, K_1 and K_2 are reset to their initial large values and slowly reduced to their final values using the same algorithm.

The logic behind the operation of the circuit is based on the principles given in [28]. The first principle is that in a locked state the phase error inside the bandwidth of the loop should be constant and the frequency error should be zero. This means that while the PLL is still locking the phase error will not be constant. The second principle is that in a locked state the PLL should approach the same locked state after small phase perturbations. This principle is more intrinsic to the operation of the PLL as a phase measurement device and basically states that if the PLL is locked small changes in phase will not cause the PLL to slip cycle.

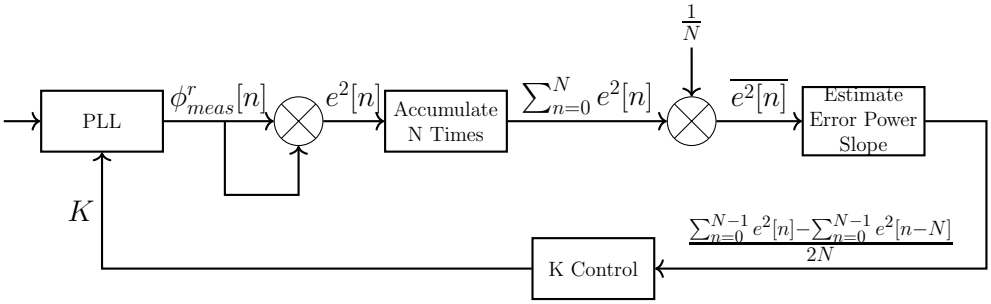

Figure 5.3 Hardware Implementation of the PLL Gain Control Algorithm

It has been found that in implementing this algorithm the lock time of the PLL is drastically reduced with no change to the operation of the circuit after it has converged. A block diagram illustrating the flow of the algorithm is shown in Figure 5.3.

5.1.2 Notch Filter

The final component in $R_1(z)$ that needs to be specified is the loop filter, $H_{lf}(z)$. This filter is not strictly necessary, but if the double frequency component arising from the multiplier phase detector needs to be suppressed, as it does for the purposes of the measurements to be made, then it becomes quite useful. The most direct and resource efficient way of removing this double frequency component is to construct a notch filter with the notch set at the frequency of the double frequency component. A pole-zero plot is shown in Figure 5.4 showing the poles and zeros of the notch filter, the zero being located at the frequency that needs to be removed on the unit circle and the pole located at the same angle, but away from the unit circle to maintain a flat pass band on the filter.

The transfer function of this filter is as follows:

$$\begin{aligned} H_{notch}(e^{jw}) &= \frac{(1 - e^{-j\phi}e^{-jw})(1 - e^{j\phi}e^{-jw})}{(1 - \rho e^{-j\phi}e^{-jw})(1 - \rho e^{j\phi}e^{-jw})} \\ &= \frac{1 - 2\cos(\phi)e^{-jw} + e^{-jw2}}{1 - 2\rho\cos(\phi)e^{-jw} + \rho^2 e^{-jw2}}, \end{aligned} \tag{5.2}$$

where ϕ is the angular frequency that needs to be removed and ρ is the parameter controlling

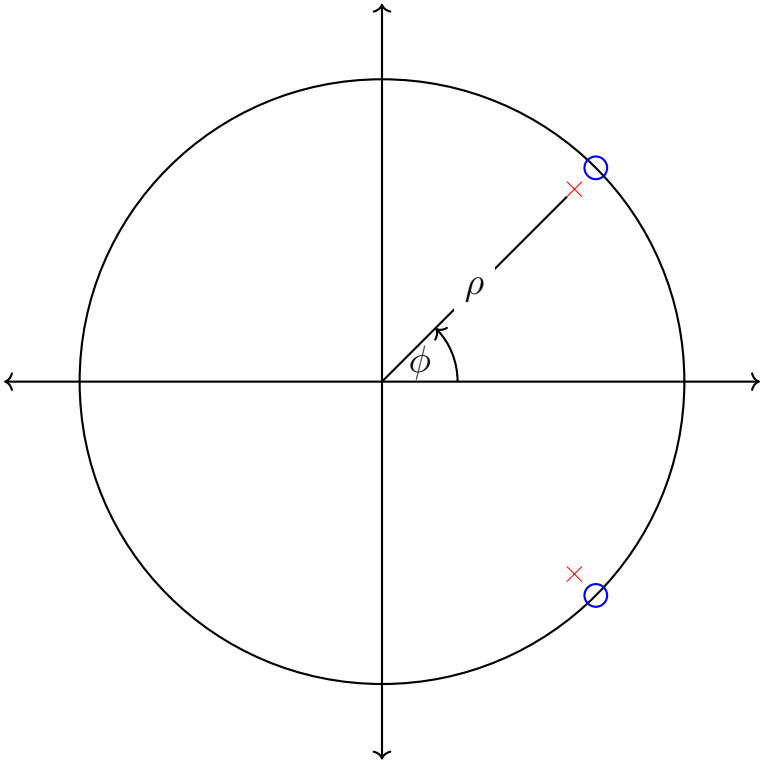

Figure 5.4 Pole Zero Plot of a Notch Filter

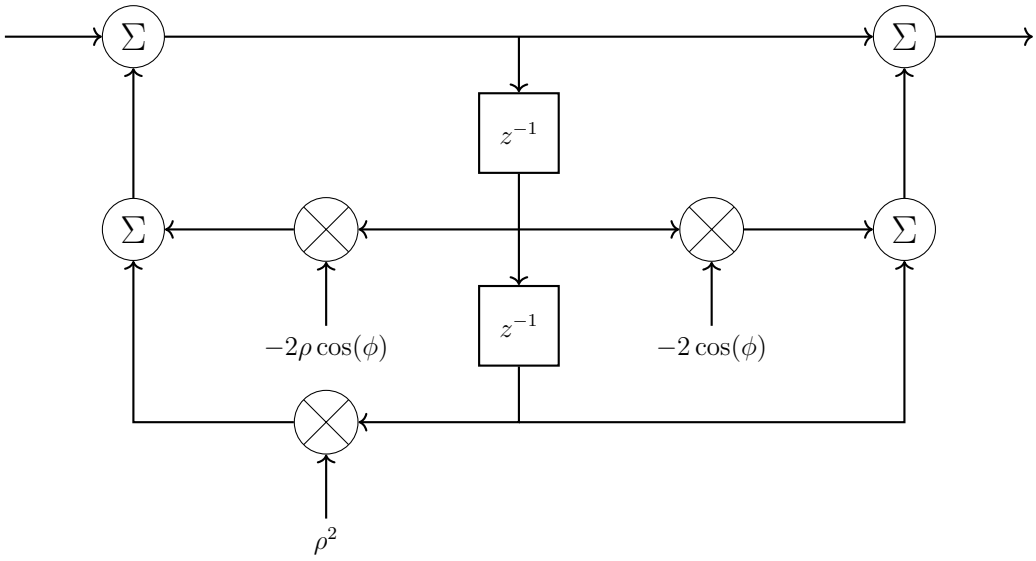

Figure 5.5 Circuit Diagram of a Notch Filter

how close to the unit circle the poles of the filter are located. A ρ closer to one yields a narrower notch with the 3dB width of the notch being equal to $2(1-\rho)$ in radians. The 3dB bandwidth of the filter must be set wide enough to allow for room in quantization error in the frequency the filter is tuned to remove. As the frequency tuning word used for the PLL is 32 bits wide, the error equals to $\Delta_{resolution} = (\frac{1}{2})^{32} \cdot F_s = 0.011$ cycles/sample with $F_s = 48$ Mhz. This makes a 3dB bandwidth of BW $= 2(1 - 0.9) = 0.2$ radians acceptable as the notch will be wide enough to accommodate the error in frequency. With a bandwidth this narrow, it will not contribute significant phase inside or near the PLLs bandwidth and will not effect the shape of the PLLs response. Including this transfer function in the feedback path of the PLL has no significant effect on the operation of the overall transfer function other than suppressing the frequency that the notch filter is set to. The circuit diagram of a notch filter realized in hardware is shown in Figure 5.5. This circuit is needed for finding the quantization noise introduced by the filter in the following section.

The parameters $2\rho\cos(\phi)$ and $2\cos(\phi)$ must be calculated based on the frequency of the NCO. The ϕ is determined based on f_0 and the notch filter coefficients are calculated as $2\rho\cos(2\pi \cdot 2f_0)$ and $2\cos(2\pi \cdot 2f_0)$. A block diagram showing how these parameters are

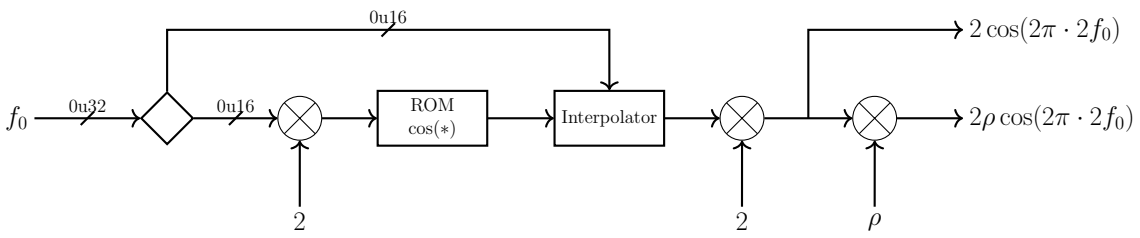

Figure 5.6 Notch Filter Coefficient Calculation

calculated is given in Figure 5.6. The parameter calculation circuit works by storing a cosine inside a ROM and using the frequency the NCO is set to as the address for that ROM. The frequency input is truncated to 16 bits to minimize the size of the ROM and the output is interpolated to increase the effective size of the ROM. From there it is scaled by the precalculated value of ρ and multiplied by 2 to get the coefficients of interest. With this circuit, the PLL will operate the same way for any frequency input, and the frequency input can be changed during operation of the circuit.

5.2 Noise Sources in the Phase Locked Loop

It is necessary to identify and quantify the noise sources inherent to the system to design the measurement circuit. This includes the sources of phase noise, thermal noise, and particularly quantization noise. After quantizing the sources of noise their effect on the measurement accuracy is analyzed.

5.2.1 Effect of Amplitude Noise on Phase Error Output

Noise that adds to the signal, like thermal noise, while additive, is not the main source of amplitude noise. In this case the amplitude noise is manifested by a channel response that changes with time thereby changing the amplitude of the received signal overtime. As such, some amount of amplitude noise is certain to be on the reference sinusoid which is the input to the PLL. Earlier on in this thesis the effect of amplitude noise in the channel was deemed insignificant when compared to the effect of the phase noise. That being said, it is still present and the multiplier phase detector converts amplitude noise to phase noise, which

then effects the phase error measurement. To determine what effect this has, the equation for the output of the phase detector is recalculated with amplitude noise on the reference sinusoid:

$$\begin{aligned} 2\pi\phi_{det}[n] &= (A_{ref} + \boldsymbol{a}[n])\cos(2\pi f_0 n + 2\pi\boldsymbol{\Delta\phi}[n]) \cdot \sin(2\pi f_0 n) \\ &= \frac{A_{ref}}{2}\sin(2\pi\boldsymbol{\Delta\phi}[n]) + \frac{A_{ref}}{2}\sin(4\pi f_0 n + \boldsymbol{\Delta\phi}[n]) + \\ &\frac{\boldsymbol{a}[n]}{2}\sin(2\pi\boldsymbol{\Delta\phi}[n]) + \frac{\boldsymbol{a}[n]}{2}\sin(4\pi f_0 n + \boldsymbol{\Delta\phi}[n]) \end{aligned} \tag{5.3}$$

Assuming the double frequency term is removed by filtering,

$$\begin{aligned} 2\pi\phi_{det}[n] &\approx \frac{A_{ref}}{2}\sin(2\pi\boldsymbol{\Delta\phi}[n]) + \frac{\boldsymbol{a}[n]}{2}\sin(2\pi\boldsymbol{\Delta\phi}[n]) \\ &\approx A_{ref}\pi\boldsymbol{\Delta\phi}[n] + \boldsymbol{a}[n]\pi\boldsymbol{\Delta\phi}[n]. \\ &\approx A_{ref}\pi\boldsymbol{\Delta\phi}[n](1 + \frac{\boldsymbol{a}[n]}{A_{ref}}) \end{aligned} \tag{5.4}$$

The noise in the amplitude of the received signal appears as $\boldsymbol{a[n]}\pi\boldsymbol{\Delta\phi[n]}$ added to the output of the phase detector. This noise is $\frac{\boldsymbol{a[n]}}{A_{ref}}$ times the phase noise. The amplitude noise is certain to be much smaller than A_{ref}, making its contribution to the noise at the output of the phase detector negligible.

5.2.2 Quantization Noise

A major source of noise in all digital systems is the noise resulting from quantizing continuous values to a number that can be represented in a finite number of bits. The first step in characterizing this noise is to examine the error resulting from quantizating an infinitely precise analog value to a finite precision quantized value.

The plot in Figure 5.7 shows an ideal input and output relation of a quantizer as the straight blue line and the quantized approximation as the red staircase with the step size of the stair case shown as s. The quantization error is defined as the difference between these twos lines and is shown as the black sawtooth. Clearly, the quantization error is bounded

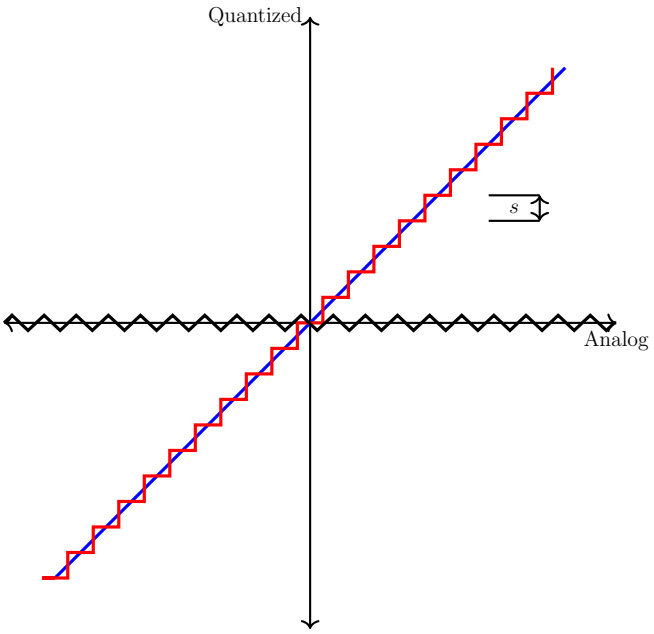

Figure 5.7 Quantized vs Unquantized

between $\pm\frac{s}{2}$. Assuming that the input is uniformly distributed, the quantization error will have a PDF (Probability Density Function) as shown in Figure 5.8 [29].

The PDF of the quantization error is modelled as a uniform distribution with zero mean. Utilizing this model the variance of the quantization error is calculated as follows:

$$\sigma^2_{qnoise} = E[(\mathbf{e}[n] - E[\mathbf{e}[n]])^2] \tag{5.5}$$

As the quantization error has zero mean the second term in the equation goes to zero and we are left with

$$\begin{aligned} \sigma^2_{qnoise} &= E[\mathbf{e}[n]^2] \\ &= \int_{\frac{-s}{2}}^{\frac{s}{2}} \mathbf{e}^2 f_e(\mathbf{e}) de \\ &= \frac{s^2}{12}. \end{aligned} \tag{5.6}$$

Where $\mathbf{e}[n]$ is the quantization error, s is the step size of the quantizer, and $f_e(\mathbf{e})$ is the

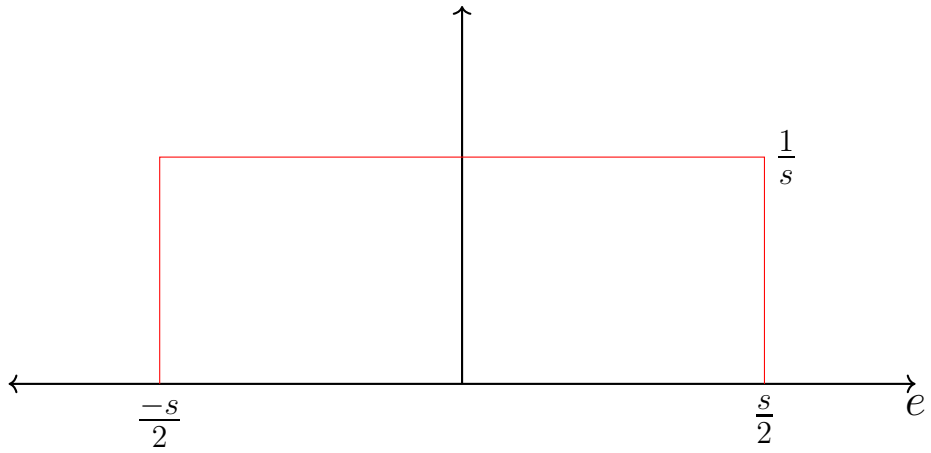

Figure 5.8 Probability Density Function of Quantization Error

PDF of the quantization error. From (5.6), the variance in the noise caused by quantizing a infinite precision analog value to a finite precision digital word is the square of the value of the least significant bit, which represents the step size, in the finite precision word divided by 12. Without loss of generality, the quantization noise is modeled throughout the rest of this thesis as white additive noise with a known variance and mean [30]. The method for calculating the mean and variance of the quantization error when moving from one word format to another is also worth knowing, the mean, μ, is given by

$$\mu = \frac{-s}{2}(1 - (\frac{1}{2})^k) \tag{5.7}$$

and the variance, σ^2, is given by

$$\sigma^2 = \frac{s^2}{12}(1 - (\frac{1}{4})^k). \tag{5.8}$$

In these formulae k is the number of bits removed and s is the value of the least significant bit. It is worth noting that (5.7) and (5.8) can also be used for an analog signal converted to a fixed precision signal. Since an analog signal has infinite precision, k is set to infinity and the formulae reduce to (5.6). This information is now used to design the lengths of the words used to implement the PLL.

Numerically Controlled Oscillator Design

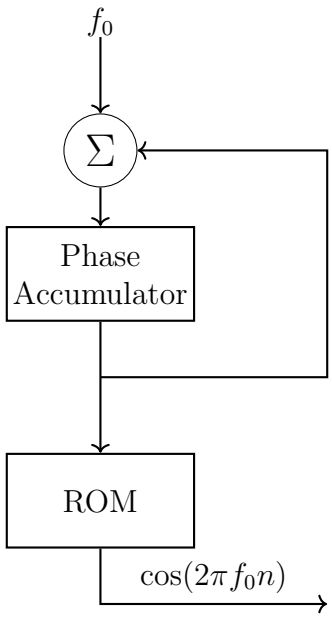

Figure 5.9 Block Diagram of a NCO

Inherent to the PLL is a feedback path that includes a Numerically Controlled Oscillator (NCO), shown in Figure 5.9. The PLL controls the input, which is the frequency, of this NCO to lock it in quadrature with the incoming sinusoid. Unfortunately, the output of a numerically controlled oscillator contains noise that is introduced due to quantization at the input and output of the ROM.

The ROM stores a single cycle of a cosine (or sine) wave [31]. The size of the words being stored in the ROM as well as the number of words being stored in the ROM determine the amount of quantization noise on the output of the NCO. The NCO must be analyzed and a metric developed to describe the noise introduced by its quantization to ensure that the noise introduced by the NCO does not significantly degrade the performance of the PLL.

The noise introduced by the NCO is modeled as two sources: truncation noise added to the data output of the ROM and truncation noise added to the address input to the ROM. The former is referred to as amplitude noise and the latter as phase noise.

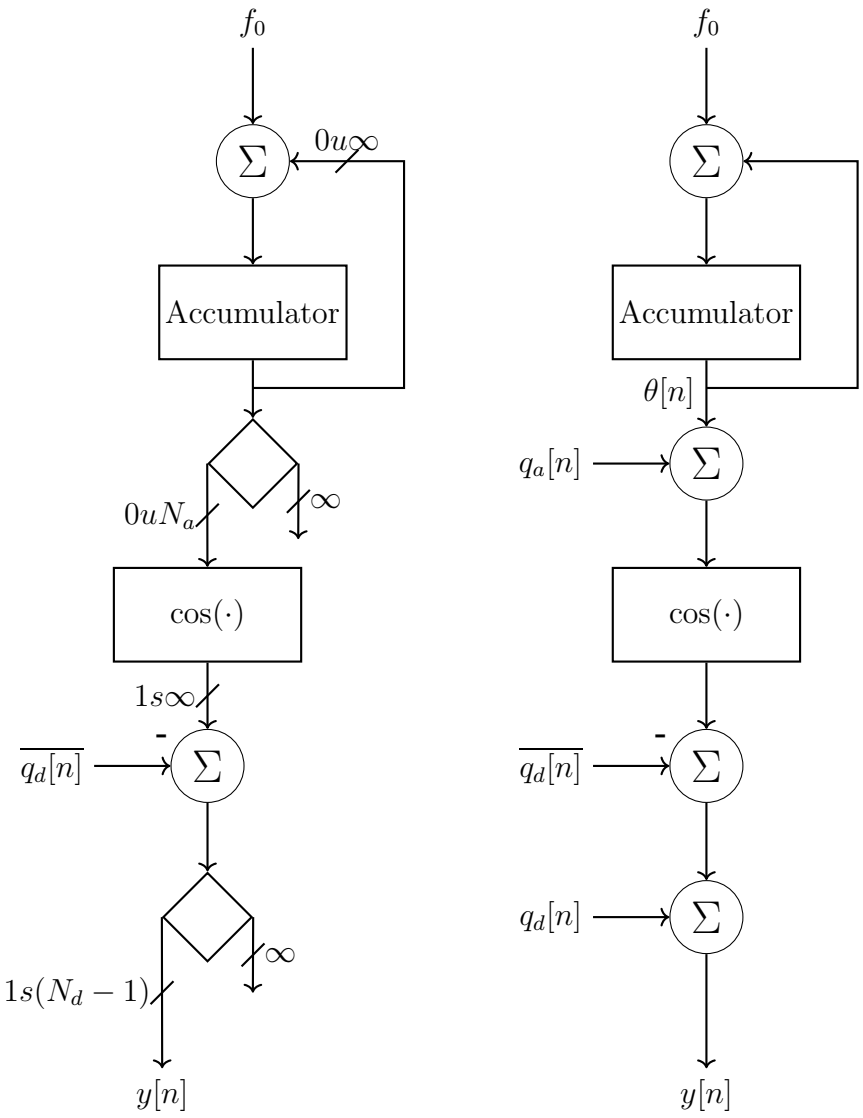

Figure 5.10 Truncated NCO Showing Bit Widths and Additive Noise

The output of the NCO circuits shown in Figure 5.10 is:

$$y[n] = \cos(\theta[n] + q_a[n]) + q_d[n] - \overline{q_d[n]} \tag{5.9}$$

where $\theta[n] = f_0 n$ with f_0 being the frequency of the cosine, $q_a[n]$ is the quantization noise due to quantizing the input to the ROM and $q_d[n] - \overline{q_d[n]}$ is the noise caused by rounding the cosine to fit in the word size of the ROM. The word size rounding noise, which is the

AC component of the truncation noise, can be calculated from (5.8) by

$$\begin{aligned} P_d &= \frac{s^2}{12} \\ &= \frac{1}{12}((\frac{1}{2})^{N_d-1})^2 \\ &= \frac{1}{3}(\frac{1}{2})^{2N_d}. \end{aligned} \tag{5.10}$$

where N_d is the number of bits in each word stored in the ROM. The noise resulting from quantizing the phase of the cosine can be converted to additive noise, but doing so is not as straight forward. Small angle approximations are utilized to calculate the additive equivalent of the phase noise. The phase noise corrupted sinusoid can be approximated by

$$\cos(2\pi\theta[n] + \mu_a + 2\pi(q_a[n] - \overline{q_a[n]})) \approx cos(2\pi\theta[n] + \mu_a) - 2\pi(q_a[n] - \overline{q_a[n]})\sin(2\pi\theta[n] + \mu_a). \tag{5.11}$$

The additive equivalent of the phase noise is the second term on the right side of (5.11). This additive equivalent of the phase noise is defined as

$$q_{\text{a\_additive}}[n] = -2\pi(q_a[n] - \overline{q_a[n]})\sin(2\pi\theta[n] + \mu_a). \tag{5.12}$$

This noise term is white because $2\pi(q_a[n] - \overline{q_a[n]})$ was found to be white by Bennett [30] and the multiplication of these two terms in the discrete time domain is equivalent to convolution in the frequency domain where a sine is an impulse. The convolution of a constant (white) term with an impulse is a constant term and thus the additive equivalent of the phase noise is white noise. The AC power of this signal is given by

$$\begin{aligned} P_a &= \overline{(2\pi(q_a[n] - \overline{q_a[n]})\sin(2\pi\theta[n] + \mu_a))^2} \\ &= 4\pi^2 \cdot \overline{(q_a[n] - \overline{q_a[n]})^2 \cdot \sin^2(2\pi\theta[n] + \mu_a)}. \end{aligned} \tag{5.13}$$

Since $q_a[n] - \overline{q_a[n]}$ is statistically independent of $\sin(2\pi\theta[n] + \mu_a)$, the time average of the product can be changed to the product of the time averages to get

$$
\begin{aligned}
P_a &= 4\pi^2 \cdot \overline{(q_a[n] - \overline{q_a[n]})^2} \cdot \overline{\sin^2(2\pi\theta[n] + \mu_a)} \\
&= 4\pi^2 \cdot \frac{1}{12}(\frac{1}{2})^{2N_a} \cdot \frac{1}{2} \\
&= \frac{\pi^2}{6}(\frac{1}{2})^{2N_a},
\end{aligned}
\tag{5.14}
$$

where $\mu_a = 2\pi\overline{q_a[n]}$ is a constant of units radians that is removed by the action of the PLL and therefore its value is inconsequential.

The total power in the quantization noise on the output of the NCO, $y[n]$, is the sum of the these two powers

$$
\begin{aligned}
P_{qn} &= P_a + P_d \\
&= \frac{\pi^2}{6}(\frac{1}{2})^{2N_a} + \frac{1}{3}(\frac{1}{2})^{2N_d}.
\end{aligned}
\tag{5.15}
$$

The power in the ideal pure cosine wave is simply $P_s = 1/2$ so the SNR at the output of the NCO can be found

$$
\begin{aligned}
\text{SNR} &= \frac{P_s}{P_a + P_d} \\
&= \frac{1/2}{\frac{\pi^2}{6}(\frac{1}{2})^{2N_a} + \frac{1}{3}(\frac{1}{2})^{2N_d}}.
\end{aligned}
\tag{5.16}
$$

From (5.16) the word lengths of the phase accumulator and words stored inside the ROM can be determined for any desired SNR.

Defining the length of the frequency input of the NCO as N_r it is possible to determine what frequency resolution you get for a specific word length. Assuming that the frequency input is a $0sN_r$ format the following equation determines what size N_r has to be for a desired frequency resolution

$$
\begin{aligned}
(\frac{1}{2})^{N_r} \cdot F_s &\leq \Delta_{\text{resolution}} \\
N_r &\geq \frac{log(\Delta_{\text{resolution}}/F_s)}{-log(2)},
\end{aligned}
\tag{5.17}
$$

where $\Delta_{\text{resolution}}$ is the desired frequency accuracy, F_s is the sampling rate, and N_r is the number of bits in the frequency input. This equation can be used to find the minimum length of N_r that achieves the specified resolution. Using (5.17), it was found that if a frequency resolution to within 0.1 Hz is required, N_r would need at least 29 bits.

From the above analysis there are three parameters that must be chosen to control the amount of noise introduced by the NCO inherent to the PLL. These parameters are the word length of the frequency input, N_r, the word length of the phase accumulator, N_a, and the word length of the output of the NCO, N_d.

Word Lengths

There are a number of quantization noise sources present in the implementation of the PLL on an FPGA. As described in the previous subsections these noise sources can be modelled as additive white noise with means and variances determined from the number of bits in the digital words used to express the signal values. The major quantization noise sources present in the FPGA implementation of the PLL measurement circuit are shown in Figure 5.12, these noise sources arise from the bit trimmers shown in Figure 5.11. The effect of each of these noise sources on the output of the circuit is determined in this subsection.

The additive noise sources shown in Figure 5.12 are as follows: $q_{adc}[n]$ is the noise introduced by the ADC, $q_{pd}[n]$ is the noise introduced by truncation after the phase detector and $q_p[n]$ the noise introduced by truncation in the pole section of the loop filter, $q_z[n]$ is the noise introduced by truncation in the zero section of the loop filter, and finally $q_{\text{a\_additive}}[n]$ and $q_d[n]$ are the noise introduced by truncating the address signal and word size of the rom discussed in the NCO section. Based on the diagram in Figure 5.12 it is possible to determine the effect that each of these noise sources has on the output of the circuit for given word sizes.

The methodology used to determine the effect that the quantization noise sources has on the desired output of the PLL is as follows. Each of the quantization noise sources shown in

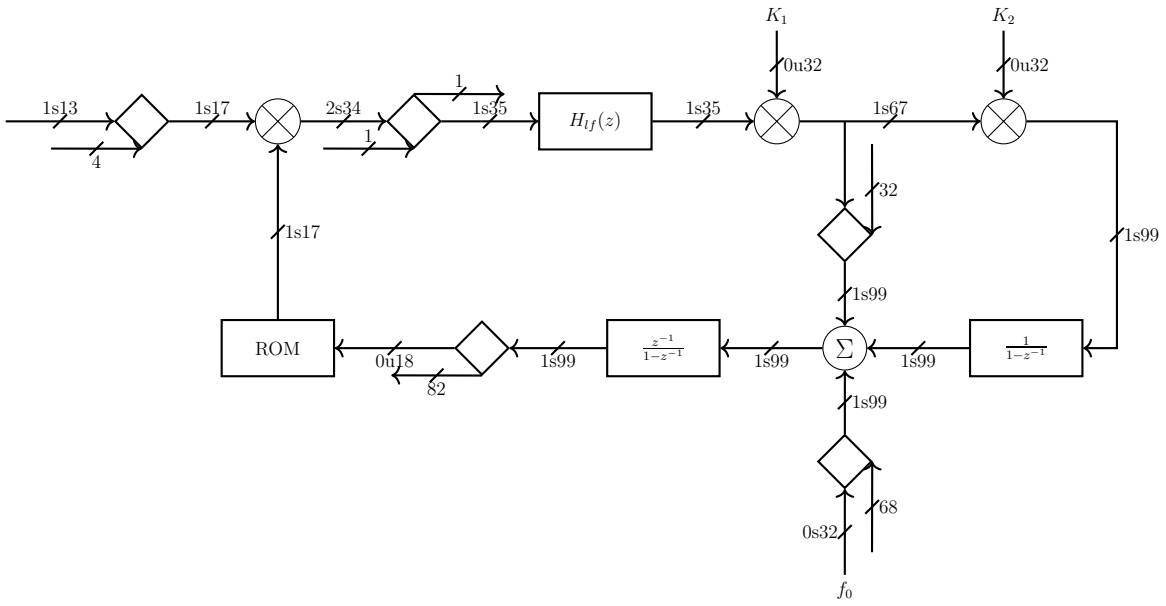

Figure 5.11 PLL Circuit Diagram Showing Word Formats

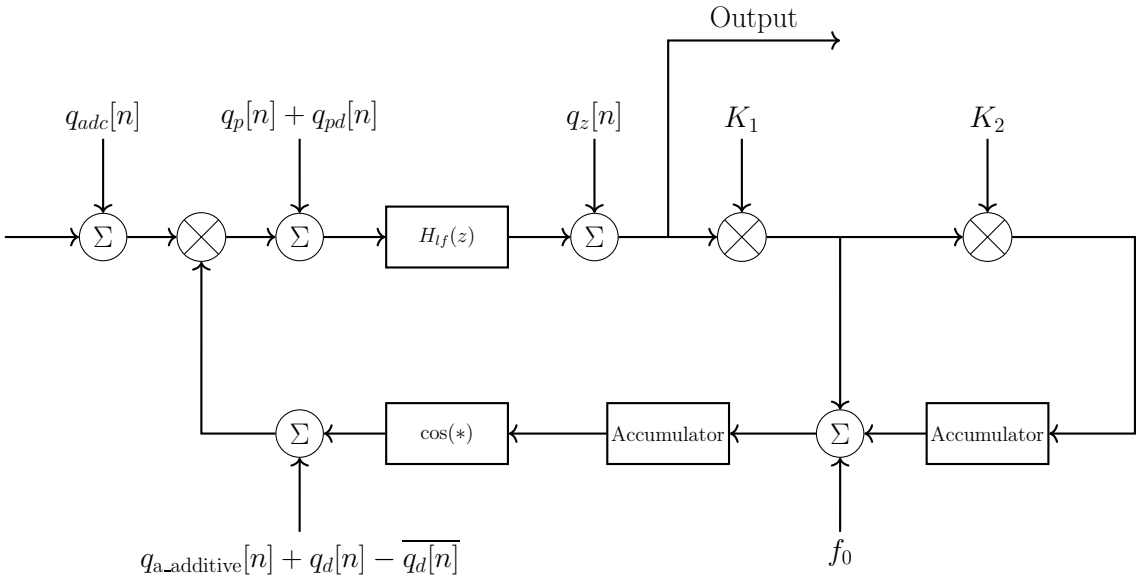

Figure 5.12 PLL Circuit with Additive Quantization Noise Sources Shown

the loop is modeled as a white noise source with a mean and variance calculated based on the word size used on the FPGA. All of the inputs besides the noise input being analyzed are set to zero and the transfer function from this input to the output of the loop is found. The power spectral density of the input is then used in tandem with the transfer function

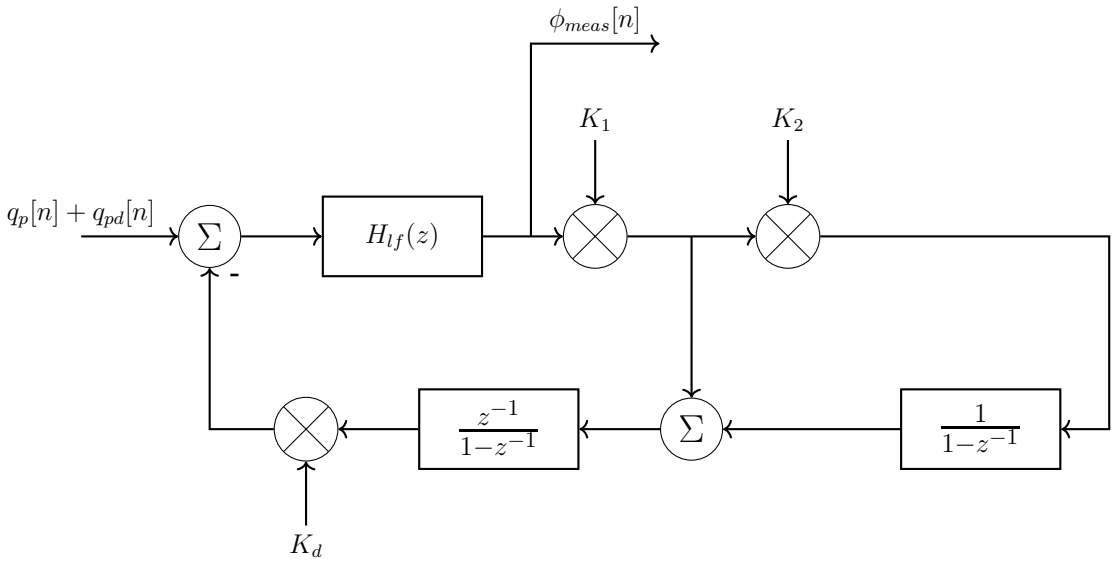

Figure 5.13 Equivalent Circuit for Quantization Noise from Phase Detector and Pole Section of Notch Filter

to determine the noise floor at the output of the circuit. The analysis makes the crucial assumption that none of the noise sources pushes the multiplier based phase detector out of its linear operation zone. If this was the case a linear analysis of the PLL could not be performed.

Since the transfer function for the ADC noise to the output is equal to the transfer function from the reference sinusoid to the output, the first step in the noise analysis is finding the transfer function from the noise introduced by the phase detector truncation and pole section of the loop filter. The equivalent circuit used to find this is shown in Figure 5.13. The open loop transfer function and the transfer function for the quantization noise to the output is given in (5.18).

$$\begin{aligned} G_{q1}(z) &= \frac{H_{lf}(z)K_1K_dz^{-1}(1-z^{-1}+K_2)}{(1-z^{-1})^2} \\ E_{q1}(z) &= \frac{H_{lf}(z)}{1+G(z)} = \frac{H_{lf}(z)(1-z^{-1})^2}{(1-z^{-1})^2 + H_{lf}(z)K_1K_dz^{-1}(1-z^{-1}+K_2)} \end{aligned} \tag{5.18}$$

The next transfer function analyzed is the transfer function for the noise originating in

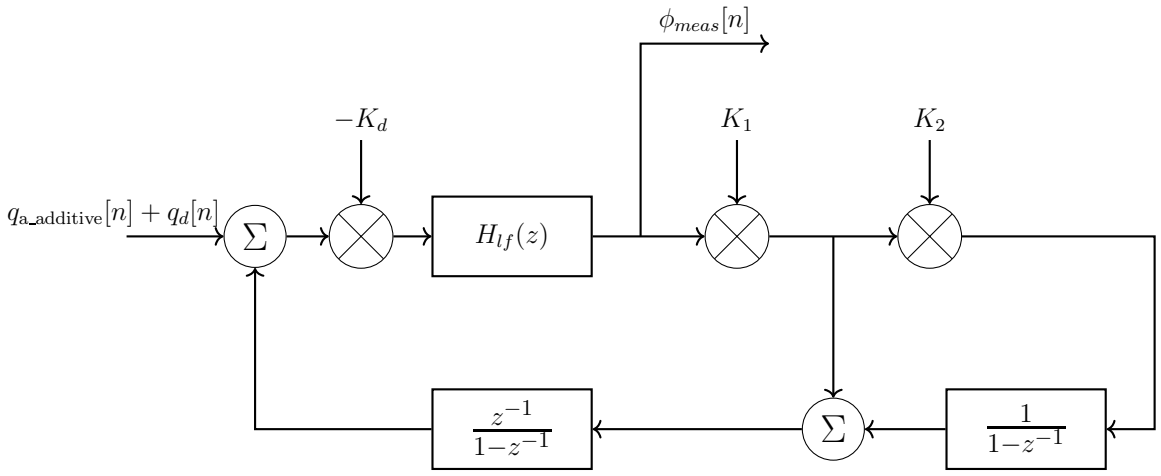

Figure 5.14 Equivalent Circuit for Quantization Noise from NCO

the NCO to the output of the circuit. The equivalent block diagram is shown in Figure 5.14. The transfer functions are shown in (5.19). It can be seen that these transfer functions are functionally identical to the previously found transfer functions, differing only in the location of the detector gain.

$$\begin{aligned} G_{q2}(z) &= \frac{-H_{lf}(z)K_1K_dz^{-1}(1-z^{-1}+K_2)}{(1-z^{-1})^2} \\ E_{q2}(z) &= \frac{-K_dH_{lf}(z)}{1-G(z)} = \frac{-K_dH_{lf}(z)(1-z^{-1})^2}{(1-z^{-1})^2-(-H_{lf}(z)K_1K_dz^{-1}(1-z^{-1}+K_2))}. \end{aligned} \tag{5.19}$$

The final transfer function needed to be analyzed is the transfer functions from the zero section of the notch filters quantization noise to the output of the circuit. The equivalent circuit is shown in Figure 5.15. The open and closed loop transfer functions are shown in (5.20).

$$\begin{aligned} G_{q3}(z) &= \frac{H_{lf}(z)K_1K_dz^{-1}(1-z^{-1}+K_2)}{(1-z^{-1})^2} \\ E_{q3}(z) &= \frac{1}{1+G(z)} = \frac{(1-z^{-1})^2}{(1-z^{-1})^2+H_{lf}(z)K_1K_dz^{-1}(1-z^{-1}+K_2)} \end{aligned} \tag{5.20}$$

Using these transfer functions the output of the system can be found as the superposition

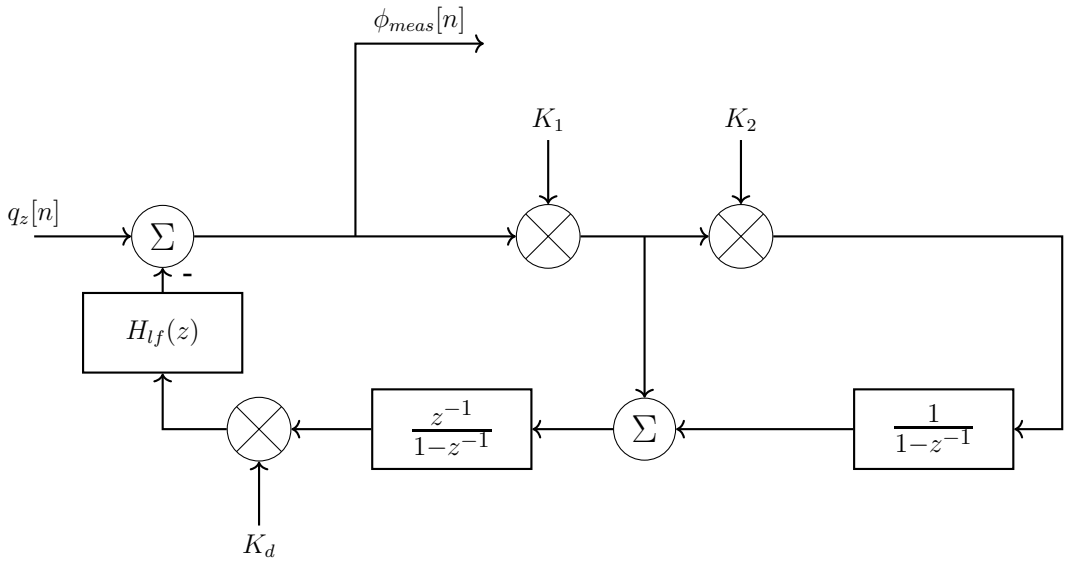

Figure 5.15 Equivalent Circuit for Quantization Noise from Zero Section of Notch Filter

of all the noise sources scaled by their respective transfer functions and the input of the system scaled by its respective transfer function. The resulting equation is given in (5.21)

$$\begin{aligned} Y(z) = & X(z)E(z) \\ & + Q_{adc}(z)E(z) \\ & + (Q_p(z) + Q_{pd}(z))E_{q1}(z) \\ & + (Q_{\text{a\_additive}}(z) + Q_d(z))E_{q2}(z) \\ & + Q_z(z)E_{q3}(z). \end{aligned} \tag{5.21}$$

In this equation the first term can be considered as the signal of interest and all other terms as noise.

The noise terms are all modeled as additive white noise with a mean and variance calculated using (5.7) and (5.8). The power spectral densities of white noise in the discrete time domain is constant across the band at a value equal to the variance of the random variable [29]

Table 5.1 Statistical Properties of Quantization Noise Sources

	s	k	σ^2 (V^2)	μ (V)
$q_{adc}[n]$	2^{-13}	∞	$2.4223 \cdot 10^{-7}$	$-6.1035 \cdot 10^{-5}$
$q_{p1}[n]$	2^{-35}	36	$7.0586 \cdot 10^{-23}$	$-1.4552 \cdot 10^{-11}$
$q_{p2}[n]$	2^{-35}	36	$7.0586 \cdot 10^{-23}$	$-1.4552 \cdot 10^{-11}$
$q_{pd}[n]$	2^{-35}	1	$5.2940 \cdot 10^{-23}$	$-7.2760 \cdot 10^{-12}$
$q_{\text{a\_additive}}[n]$	2^{-18}	∞	$2.3937 \cdot 10^{-11}$	$-1.1984 \cdot 10^{-5}$
$q_d[n]$	2^{-17}	∞	$4.8506 \cdot 10^{-12}$	0
$q_z[n]$	2^{-35}	36	$7.0586 \cdot 10^{-23}$	$-1.4552 \cdot 10^{-11}$

$$S_{qxqx}(e^{jw}) = \sigma_x^2 \, \frac{\mathrm{V}^2}{\text{cycles/sample}}. \tag{5.22}$$

The power spectral density of the incoming signal is not strictly known but to determine acceptable word lengths it is enough to know that the circuit must be able to detect signals with a power of $10^{-7}\,V^2$. The initial word lengths decided upon are summarized in Table 5.1 and shown in Figure 5.11. The values found in Table 5.1 were calculated using (2.22), (5.13), (5.10), (5.7) and (5.8). The ENOB used to calculate the ADC noise was set at 9.6 which is the value specified by the manufacturer. It is noted that even though N_a and N_d are both 18 bits long, N_a uses an unsigned format and therefore the step size is half of that used in N_d.

With the equivalent transfer functions for the different quantization noise sources inside the loop found as well as the power spectral densities of the noise sources it is possible to determine what the noise floor of the measurements will be. As the noise sources are all modeled as white noise their power spectral densities are flat across the band. The magnitude response of $H_{lf}(z)$ is assumed to be unity across the band except for the frequency that it is designed to suppress, which is sufficiently narrow so as not to affect the results. The noise transfer functions also have unity gain across the band except for $E_{q2}(z)$ which is scaled by K_d. With this information the power spectral density of the noise floor in $\frac{V^2}{\text{cycles/sample}}$ is determined by setting $K_d = A_{ref}/2 = 0.7/2$, which was found experimentally. It is noted that the DC components of the quantization noise sources will be removed by the action of the PLL and can be ignored.

$$
\begin{aligned}
Q_{noise\_floor}(z) &= S_{qadcqadc}(z)E^2(z) \\
&\quad + (S_{qpqp}(z) + S_{qpdqpd}(z))E_{q1}^2(z) \\
&\quad + (S_{qaqa}(z) + S_{qdqd}(z))E_{q2}^2(z) \\
&\quad + S_{qzqz}(z)E_{q3}^2(z) \\
&= 2.4223 \cdot 10^{-7} \cdot 1 \\
&\quad + (2 \cdot 7.0586 \cdot 10^{-23} + 5.2940 \cdot 10^{-23}) \cdot 1 \\
&\quad + (2.3937 \cdot 10^{-11} + 4.8506 \cdot 10^{-12}) \cdot (0.7/2)^2 \\
&\quad + (7.0586 \cdot 10^{-23}) \cdot 1 \\
&\approx 2.4223 \cdot 10^{-7} \, \frac{V^2}{\text{cycles/sample}}.
\end{aligned}
\tag{5.23}
$$

The noise floor at the output of the PLL measurement circuit is completely dominated by the noise from the ADC as can be seen from (5.23). This revelation implies that with the ADC used for this design, the circuit will not be able to detect a signal with variance of 10^{-7}. However, with this design the PLL can still come close to meeting this goal. With these findings the PLL circuit is ready to be implemented in hardware in Chapter 6.

6. Verification of Phase Locked Loop Measurement Circuit

This chapter deals with verifying the operation of the PLL circuit designed in the previous chapters. The theoretical operation of the circuit is first simulated in MATLAB to ensure it operates as designed. After the circuit is verified in MATLAB it is built on an Intel Arria 10 FPGA and tested to ensure it is ready to make the measurements of interest.

6.1 MATLAB Verification

The verification of the design begins by simulating the conditions of the measurement circuit in MATLAB.

6.1.1 Program Flow Overview

The MATLAB simulation can be described by the sequence of actions listed below:

1. The ROM in the NCOs is constructed in MATLAB by storing a single cycle of a cosine inside a vector. The amplitude of the cosine is first quantized before it is stored in the vector.

2. A phase accumulator for the transmit NCO is constructed in MATLAB by creating a staircase wave. The staircase starts at the initial phase set at the beginning of the simulation and ramps to 2^{N_a}, where N_a is the number of bits in the address signal of the ROM. The step size of the staircase is determined by the frequency that the generated sinusoid is required to be and is set at the beginning of the simulation.

3. A test signal is added to the staircase phase signal, which is equivalent to phase modulating the generated sinusoid. This test signal is set as a single low frequency tone with a user selected amplitude and frequency.

4. The modified staircase waveform is then used as the address for the NCO ROM, this generates a cosine wave phase modulated by the test signal. The number of samples in the simulation is equal to the length of this cosine and is controlled by a parameter at the start of the simulation.

5. The effects of the channel on the transmitted cosine are then simulated. First, the noise from the phase jitter that the DAC would introduce to the signal is modelled and then applied to the sinusoid. A number of echoes are added to the transmitted signal by adding delayed and scaled versions of the transmitted signal together. With the echoes added, the effects of a hybrid coupler are then added to the transmitted signal. The effect of the hybrid is to shift the sinusoid 90° out of phase and attenuate it by 30 dB. Finally, thermal noise is added to the transmitted signal.

6. The effects of the ADC on the received signal are then modeled. The ADC hard limits the signal between -1 and 1 as well as adding quantization noise based on the effective number of ADC bits as discussed in Chapter 2.

7. The sampling rate to this point has been 3.072 GHz, however this is far too high to be implemented on an FPGA. To address this issue the signal is down sampled in stages by a factor of 64. Before each stage of down sampling the signal is filtered using a half band filter to prevent aliasing. This brings the sampling rate of the received signal down to 48 MHz, which is the rate that it runs in hardware.

8. At this point the transmitted signal includes the effects of the channel and is ready to enter the simulated PLL. The PLL is modeled using a large for-loop that iteratively updates the vectors inside the loop. Inside the loop, the received signal is first multiplied by a locally generated sinusoid originating from an NCO inside the PLL, this

mimics the effects of the multiplier based phase detector. The output of the phase detector is then filtered with a notch filter to suppress the double frequency components. The filtered phase signal is then scaled by the specified values of K_1 and K_2. From here the output of the K_2 multiplier is sent into the frequency accumulator, the output of which is added to the output of the K_1 multiplier as well as the frequency of the transmitted sine wave to create the input of the local phase accumulator. The output of the frequency adder is then sent into the local NCO's phase accumulator which is wrapped to create another staircase waveform representing the phase of the local NCO. The staircase is then again used as the address in a cosine ROM, the output of which is multiplied by the received reference signal and the process repeats until the loop finishes. Throughout the loop the pre-calculated quantization noise is added in the positions specified in Chapter 5 in the form of white noise.

9. After the loop is finished the results are plotted and observed by the user. If the loop failed to lock in the number of iterations specified the option is given to save the current state of the vectors in the loop and run it again. If the loop did achieve lock then its output is plotted and compared to the theoretical test signal to determine if the phase variation set by the user is properly recovered.

The MATLAB simulation is summarized in the flow chart shown in Figure 6.1.

6.1.2 MATLAB Simulation Results

The MATLAB script was used to verify the theoretical operation of the design. A fairly large value of 2^{-13} was used for K_1 to illustrate the operation of the PLL. The large K_1 provided faster convergence times ensuring that the script would only need to be ran once.

The test signal used was a cosine wave at a frequency of 2500 Hz a power of $10^{-7}V^2$. The frequency of the sinusoid that was transmitted to the network was 10 MHz. This sinusoid was then phase modulated by the test signal. Noise was applied to the transmitted signal throughout the script at the positions and powers outlined at the end of Chapter 5. The

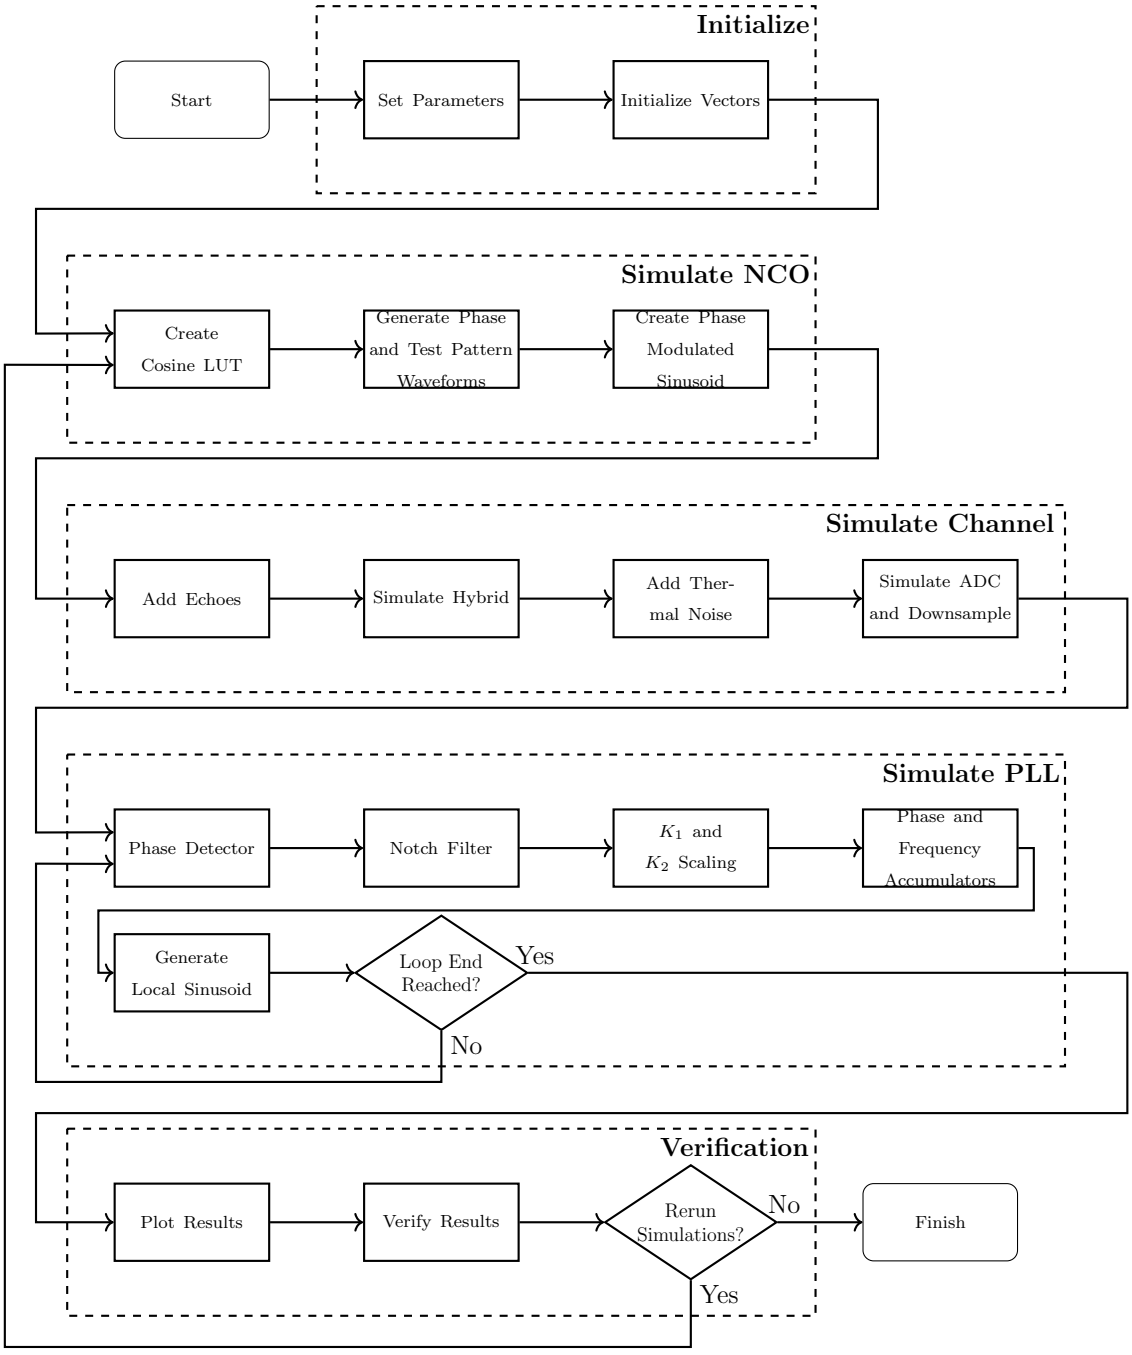

Figure 6.1 MATLAB Simulation Flow Chart

clock jitter noise added to the signal was specified in the frequency domain as having power [-81 -98 -118 -128 -124] dB below the carrier at frequencies [10 100 100 1000 10000] Hz both above and below the carrier frequency. The points between these were interpolated and

the frequency mask applied to the transmitted signal to generate the characteristic spectral spreading associated with jitter noise at the specified power and frequencies. In Chapter 5, the analysis showed that the ADC noise would swamp a signal with power $10^{-7}V^2$, to ensure that the theoretical operation of the circuit is working as expected the ENOB in the ADC was set at 12, instead of the true value of 9.6 so that the test signal could be properly recovered.

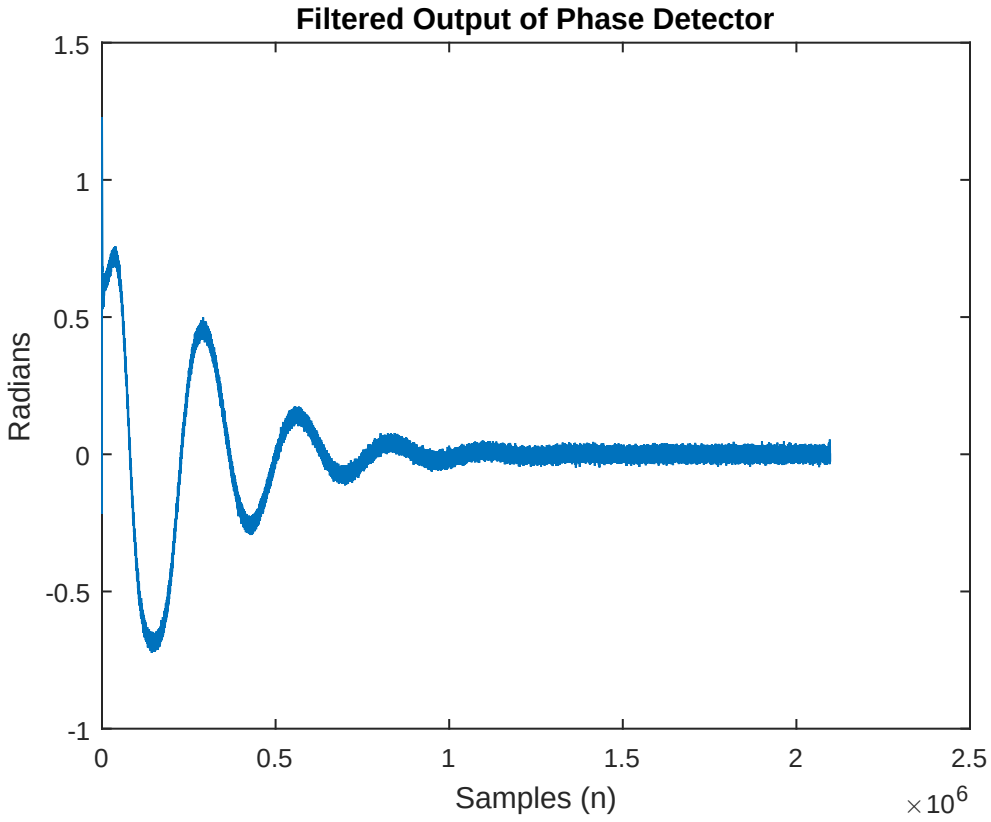

Figure 6.2 Filtered Output of Phase Detector

The results of the MATLAB script are plotted in Figure 6.2 and Figure 6.3. Figure 6.2 shows the output of the simulated phase detector after it has been passed through a notch filter. It is clear that after some initial oscillations the output converges to a steady state. In this example it takes approximately $1.5 \cdot 10^6$ samples to converge. At a sampling rate of $48 \cdot 10^6$ samples/second the convergence time is approximately 0.03125 seconds. Upon looking at the output of the phase detector after convergence the test signal is not readily

discernible from the noise. We must look at the spectral content of the waveform to confirm that the test signal is present.

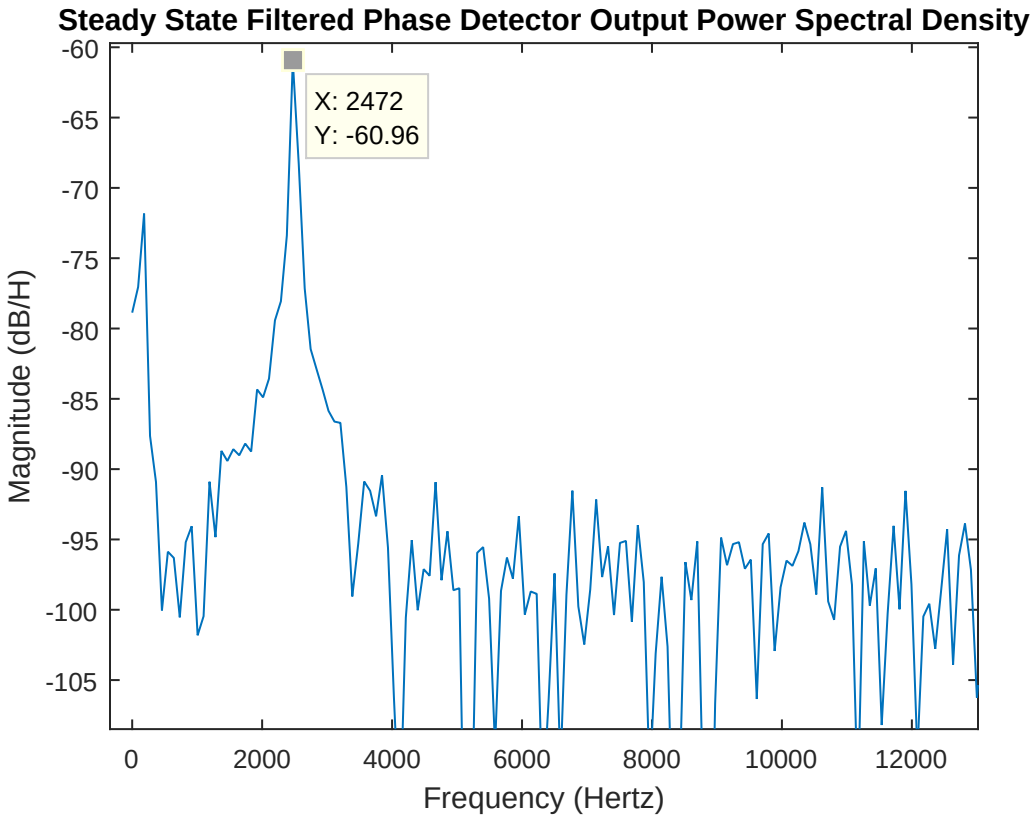

Figure 6.3 Zoomed Power Spectral Density of Filtered Phase Detector Output

Figure 6.3 shows a zoomed in power spectral density of the output of the phase detector after the phase detector has converged. Clearly there is a peak at ≈ 2500 MHz with magnitude -61 dB. The transmitted power of the test signal was -70 dB, but the detector gain of the PLL produces a level of -61 dB.

It can also be seen that there is a second peak at a low frequency. This peak is present because the loop is not truly in steady state. The "start up" oscillations were still present when the power spectral density was taken. If the data used for Figure 6.3 was collected farther into the simulation this second peak would completely vanish. From this power spectral density it becomes clear that the test signal is present in the output of the phase

detector verifying that the phase locked loop is able to detect phase fluctuations as predicted.

6.2 Verification of FPGA Implementation

Having designed the measurement circuit and verified its operation in MATLAB the design is ready to be implemented in hardware. The FPGA used to build the design is an Intel Arria 10 SoC. Connected to the FPGA test board via FMC (FPGA Mezzanine Card) is a combined Analog Devices AD9162 DAC and Analog Devices AD9208 ADC development board. A photo of the hardware components is shown in Figure 6.4. A block diagram outlying the different hardware components and subsystems is shown in Figure 6.5.

It is pointed out that designing circuits to run at clock rates as high as 3.072 GHz is very challenging. Substantial iterative effort was spent debugging and redesigning the circuit to ensure the set-up and hold requirements of the flip flops in the FPGA, as well as those of the two chips connected to the FPGA were met.

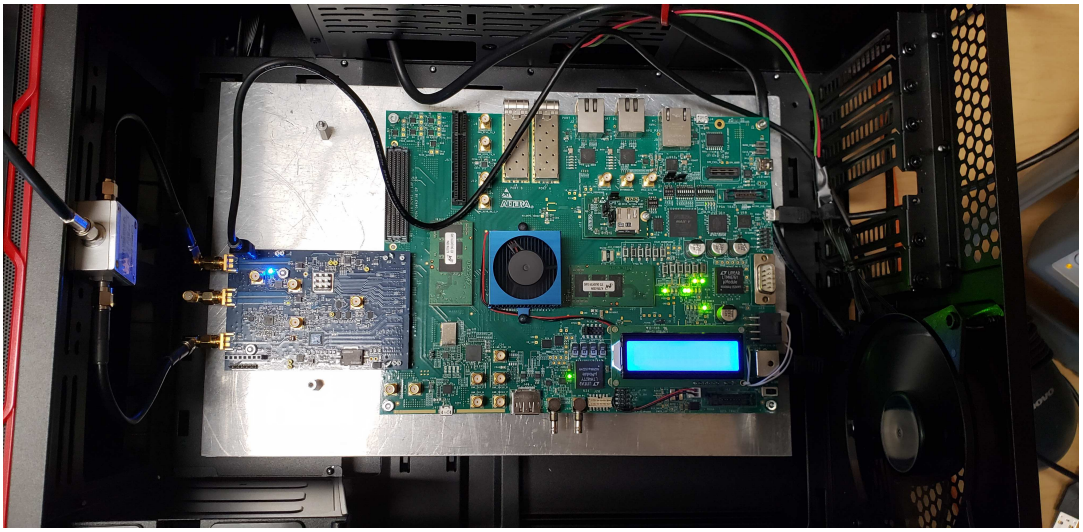

Figure 6.4 FPGA and DAC/ADC Test Boards Used for the Measurement Circuit

6.2.1 Transmit Circuitry

The connection between the FPGA and DAC is very complicated since the DAC runs at 3.072 GHz and the FPGA runs at 384 MHz. Due to the high clock rates involved, the

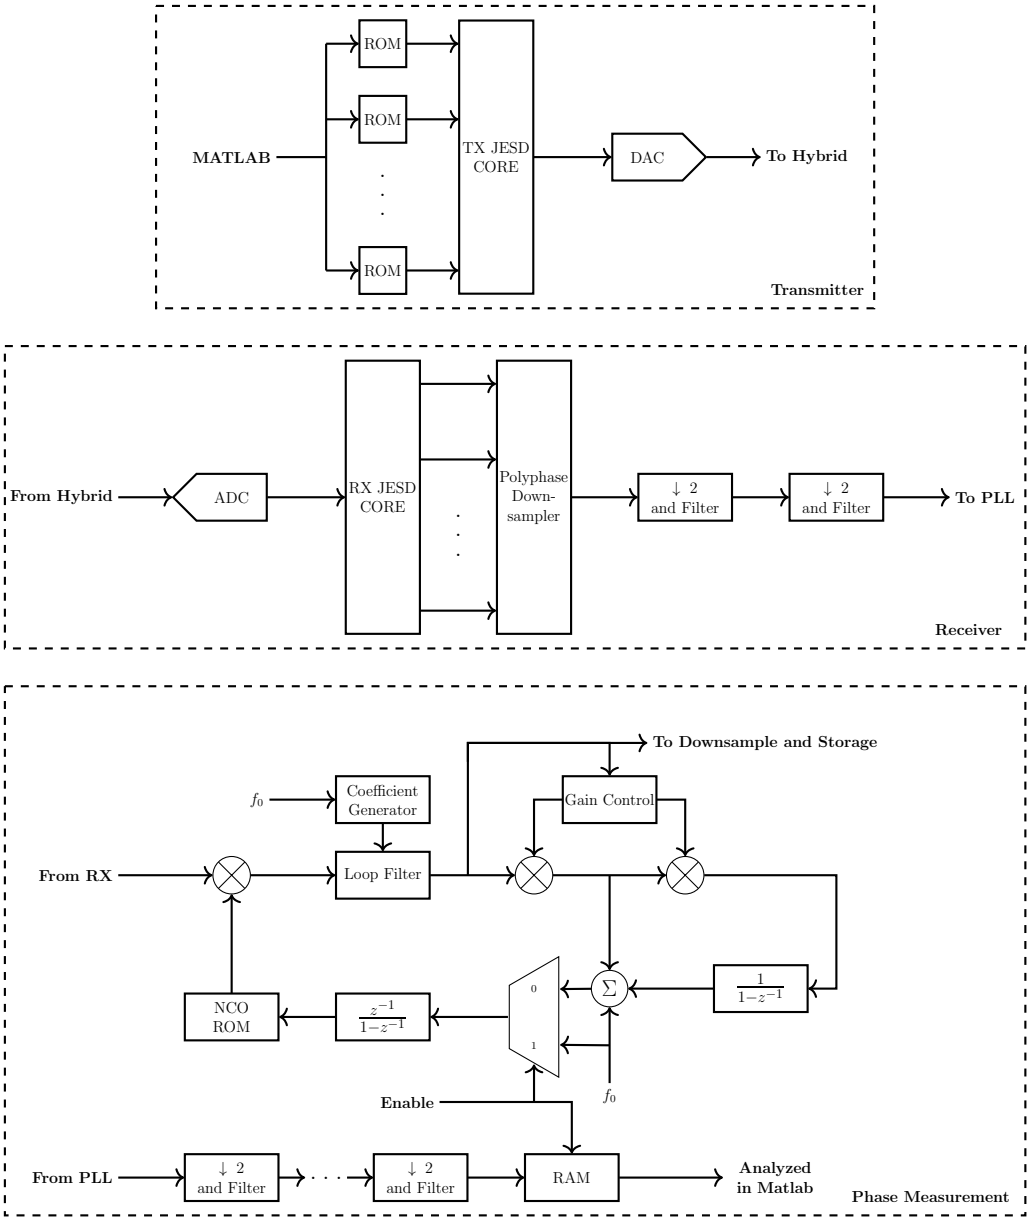

Figure 6.5 Hardware Block Diagram

connection between the FPGA and the DAC uses a JESD204b interface. The input to the interface is organized into eight parallel lanes each clocked at 384 MHz. The output of the interface to the DAC can be thought of as a single lane with a sampling rate of 3.072 GHz. The JESD204b is a high speed protocol, and it is necessary to interface the extremely high rate of the DAC with the FPGA. Implementing a serial connection between the DAC and

the FPGA at 3.072 GHz would be impossible with current technology, the eight parallel lanes at 384 MHz is much more reasonable.

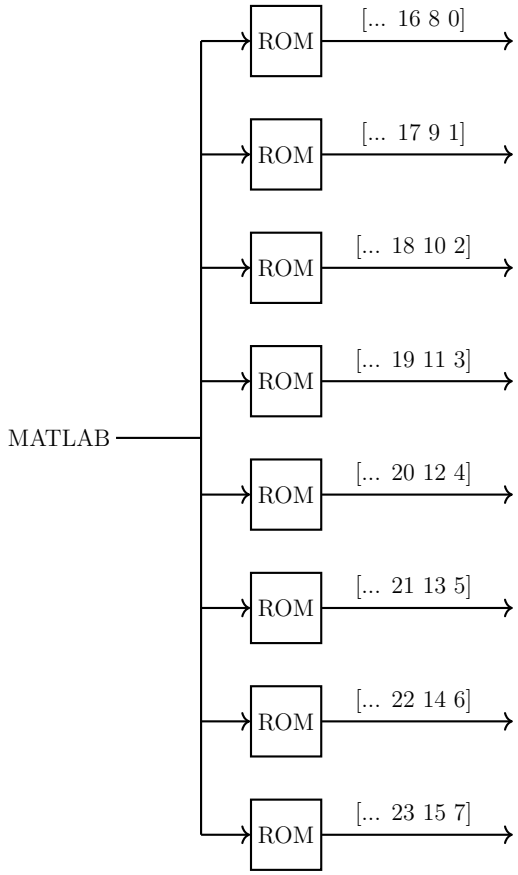

Figure 6.6 Transmitted Signal Order

Attached to each of the eight parallel lanes feeding the JESD204b interface is a ROM containing 2^{11} addresses with each address storing a 16 bit word. Each of the eight ROMs is loaded with one eighth of any arbitrary signal with each successive ROM generating the next sample in the sequence. For example, if the signal being transmitted was a sawtooth wave with a slope of 1 and a period of 16384, i.e. started at 0 and is reset every 16384 samples, the zeroth address in the zeroth ROM would contain 0, the zeroth address in the first ROM would contain 1, the zeroth address in the second ROM would contain 2, and so on and so forth. The transmission ordering process is illustrated in Figure 6.6. The ROMs are loaded with .mif files generated in MATLAB and can be updated from a PC during operation of the

circuit via JTAG allowing for the test signal to be changed without recompiling the verilog HDL.

In this system the transmitted signal is a single sinusoid at 189.75 MHz. This frequency was chosen as it is in the full duplex band, high enough that the input/output impedance of the DAC/ADC is at its specified value of 75Ω, and is equivalent to a 2.25 MHz sinusoid when sampled at 192 MHz [32]. The output of the DAC was connected via a short 50Ω cable to a 3 GHz Agilent Technologies CXA Spectrum Analyzer, a screen capture of which can be seen in Figure 6.7.

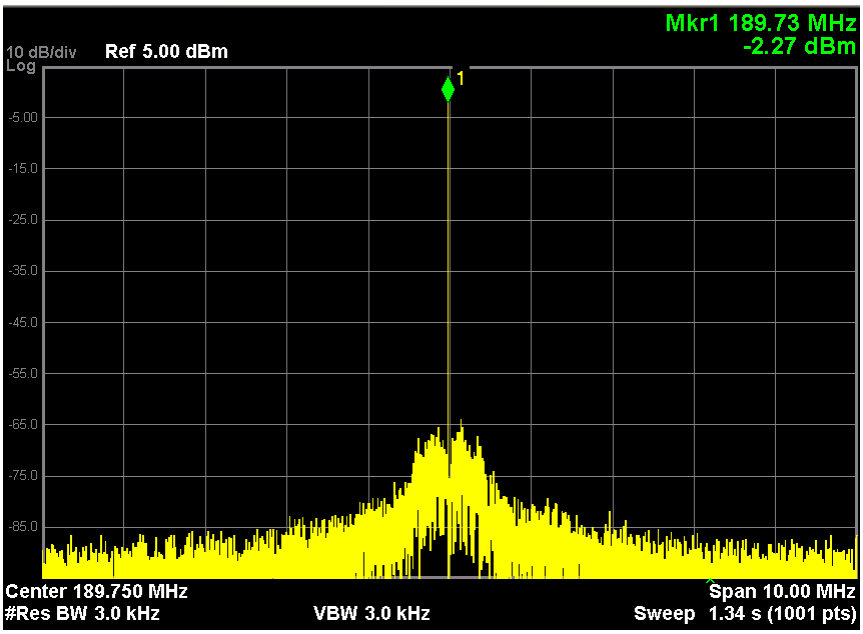

Figure 6.7 Dac Output

The transmitted sinusoid is seen to be delivering -2.27 dBm to the spectrum analyzer. The DAC is able to deliver more power to the network than this but due to the spectrum analyzers 50Ω input impedance and the DAC's 75Ω output impedance max power transfer is not achieved. It can be seen in Figure 6.7 that the transmitted sinusoid has some phase noise associated with it causing the characteristic spectral spreading. This phase noise is to be expected and was discussed at length in Chapter 3 and is caused by the clock jitter discussed in Chapter 2. This transmitted sinusoid is then sent through some arbitrary network and

received by an ADC.

6.2.2 Receive Circuitry

The transmitted signal is received by an AD9208 Analog Devices ADC. The interface between the ADC and the FPGA is again achieved using a JESD204b interface. Instead of the 8 parallel lanes clocked at 384 MHz like on the transmit side, the received signal is organized into 16 parallel lanes each clocked at 192 MHz. The equivalent sampling rate on the received side is 3.072 GHz which is far higher than needed and also very difficult to process on an FPGA. To remedy this issue, a series of downsampling stages are implemented to bring the sampling rate to a more reasonable value. The first stage of this downsampling is by a factor of 16 and uses a polyphase filter to reduce the 16 parallel lanes into 1 lane clocked at 192 MHz.

Polyphase Anti-Aliasing Filter

The polyphase filter works under the Noble Identity which states that if the output of a filter is followed by an M-to-1 downs sampler the filter can be downsampled and moved to the other side of the downsampler [33]. Application of this Noble Identity is illustrated in Figure 6.8.

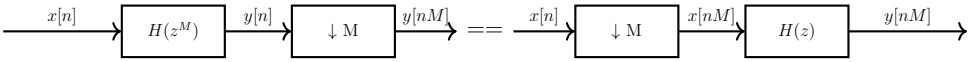

Figure 6.8 Noble Identity

Insight into the noble identity is gained by synthesizing the system function of an arbitrary filter, say $H(z)$, into M filters as follows:

$$\begin{aligned} H(z) &= \sum_{n=0}^{N-1} h[n]z^{-n} \\ &= h[0] + h[1]z^{-1} + h[2]z^{-2} + ... + h[N-1]z^{-N+1} \end{aligned} \tag{6.1}$$

The end goal is to downsample by M (in this case M = 16) so the sum shown in (6.1) is

re-partitioned into a sum of M sums.

$$\begin{array}{llll} H(z) = h[0] & + h[M+0]z^{-M} & + h[2M+0]z^{-2M} & + \ldots \\ \quad + h[1]z^{-1} & + h[M+1]z^{-M-1} & + h[2M+1]z^{-2M-1} & + \ldots \\ \quad + h[2]z^{-2} & + h[M+2]z^{-M-2} & + h[2M+2]z^{-2M-2} & + \ldots \\ \quad + h[3]z^{-3} & + h[M+3]z^{-M-3} & + h[2M+3]z^{-2M-3} & + \ldots \\ \quad \vdots & \vdots & \vdots & \vdots \\ \quad + h[M-1]z^{-M+1} & + h[2M-1]z^{-2M+1} & + h[3M-1]z^{-3M+1} & + \ldots \end{array} \tag{6.2}$$

Each row in (6.2) is itself a filter response operating on every Mth sample. The sum of the nth row in (6.2) can be expressed as $H_n(z^M)z^{-n}$ for $n = 0 : M-1$. Using this terminology the sum of sums in (6.2) can be expressed as

$$H(z) = H_0(z^M) + z^{-1}H_1(z^M) + z^{-2}H_2(z^M) + z^{-3}H_3(z^M) + ... + z^{-(M-1)}H_{M-1}(z^M). \tag{6.3}$$

This result shows that we can take an arbitrary filter and decompose it into M filters each operating on every Mth input, delay the input to the nth path by n and then sum the results and end up with the same response as the original filter. This suits the needs for down sampling perfectly as the output of the RX JESD204b core is organized into 16 parallel lanes with each lane being delayed by 1 sample with respect to the previous lane. (6.3) shows that we can filter the output of each of these lanes with the respective decomposed filter, then sum the outputs to get the filtered high sampling rate signal. This high rate signal can then be downsampled without fear of strong aliases. This process is illustrated in Figure 6.9.

The final step in the polyphase filter design is to utilize the Noble Identity to move the down sampler to the input side of the filter as shown in Figure 6.10. In doing this each filter runs at the slow rate, processes every n^{th} sample of the input, and the output is the down sampled and filtered version of the input. The input to the partitioned filters now

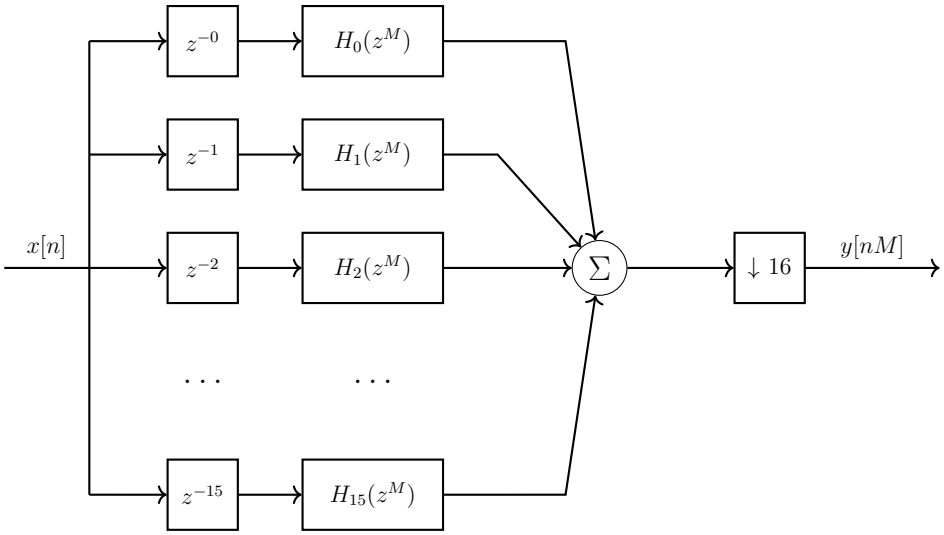

Figure 6.9 Polyphase Partition of a Low Pass Anti Aliasing Filter

exactly mimics the signal coming out of the JESD204b core, each lane is clocked at 1/16th the overall sampling rate and delayed by one with respect to the previous lane.

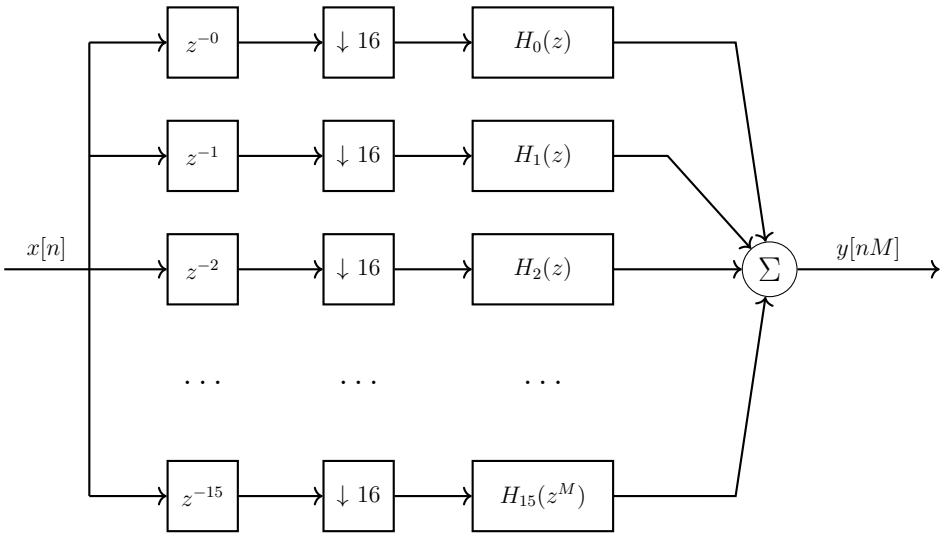

Figure 6.10 Polyphase Partition of a Low Pass Anti Aliasing Filter With Input Down Sampling

By implementing the filter as shown in Figure 6.10 and setting the cut off frequency of the filter to 1/16 cycles/sample the signal can be filtered and down sampled by 16 all while working at a sampling rate of 192 MHz.

The output of the polyphase down sampling filter is down sampled in two stages to a rate of 48 MHz. Prior to each of these down sampling stages the signal is pre-filtered to prevent aliasing with a 2-path FIR half band filter. It is noted that recursive all pass based half band filters are more resource efficient than their FIR equivalent but due to ease implementation FIR filters are used [34].

Downsampling and filtering after the polyphase filter has the effect of severely limiting the range of frequencies that the PLL is able to process. Frequencies higher than the sampling rate of $F_s = 192$ MHz are possible to measure because of deliberate aliasing from higher Nyquist zones but the subsequent anti aliasing filtering used to bring the sampling rate down to 48 MHz limits the range that the PLL is able to process. Due to this, the measurement circuit can only process frequencies located in the ranges of $nF_s \pm F_{PLL}/2$ MHz for $n = 0, 1, 2, \ldots$ where F_{PLL} is the sampling rate of the PLL. This range can be expanded by running the PLL at a higher frequency but this also has the effect of increasing the number of registers required in the PLL as well as adding more implementation complexity. If it is desired to have a wider range of measurable frequencies, one or two stages of downsampling can be removed. It is also noted that the K values in the PLL would have to be recalculated in this case. A diagram showing the frequency range of interest and the areas in that range that the PLL can measure while being ran at 48 MHz is shown in Figure 6.11.

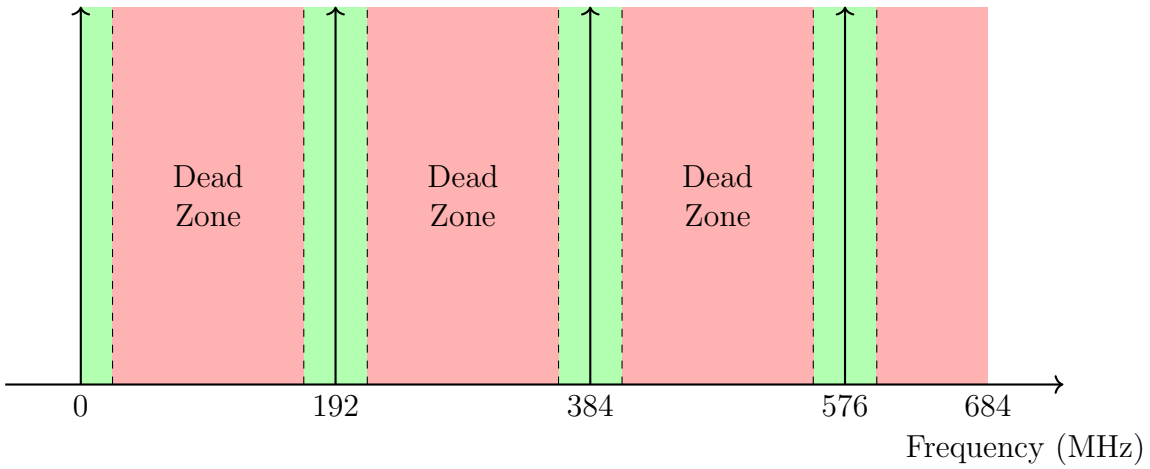

Figure 6.11 Measurable Frequency Zones of PLL Circuit

Amplitude Recovery Circuit

The gain of the phase detector depends on the amplitude of the received sinusoid therefore a circuit is required to recover the amplitude, A_{ref}, of the incoming sinusoid. This amplitude is referred to as the reference level, hence it is denoted A_{ref}. The reference level is found by multiplying the incoming signal with the in-phase output of the NCO once the NCO is locked as the in-phase output of the NCO is phase locked with the received sinusoid. Following this multiplication the double frequency component is then removed using the same notch filter specified in Chapter 5. Since the amplitude of the NCO output is 1, the DC component remaining after filtering is $A_{ref}/2$. The mathematics that supports this argument is shown in (6.4).

$$\begin{aligned} A_{recovered} &= (A_{ref}\sin(2\pi f_0 n + 2\pi\phi[n]) \cdot \sin(2\pi f_0 n + 2\pi\phi[n])) * h_{notch}[n] \\ &= (\frac{A_{ref}}{2}\cos(0) - \frac{A_{ref}}{2}\cos(4\pi f_0 n + 4\pi\phi[n])) * h_{notch}[n] \\ &\approx \frac{A_{ref}}{2}. \end{aligned} \tag{6.4}$$

The resulting value, $A_{recovered}$, is the scaling factor needed to properly measure the phase fluctuations. A block diagram illustrating the amplitude level recovery process is shown in Figure 6.12. In Figure 6.12 there are two outputs of the PLL, one term is quadrature and is sent to the phase detector and the other is in-phase and is sent to the amplitude detector.

Implemented Circuit Output

With a sampling rate of 192 MHz the 189.75 MHz sinusoid transmitted to the network aliases to a 2.25 MHz sinusoid due to its location in the second Nyquist zone. Therefore, the PLL is tuned to lock to a 2.25 MHz sinusoid and generates the phase fluctuations of the signal as described in Chapters 4 and 5. The output of the phase detector is then further down sampled down to a sampling rate of 11.71875 KHz to minimize the memory requirements of the data storage circuits. Prior to each stage of down sampling the phase fluctuations are again band limited by a FIR half band filter to prevent aliasing. This sampling rate is

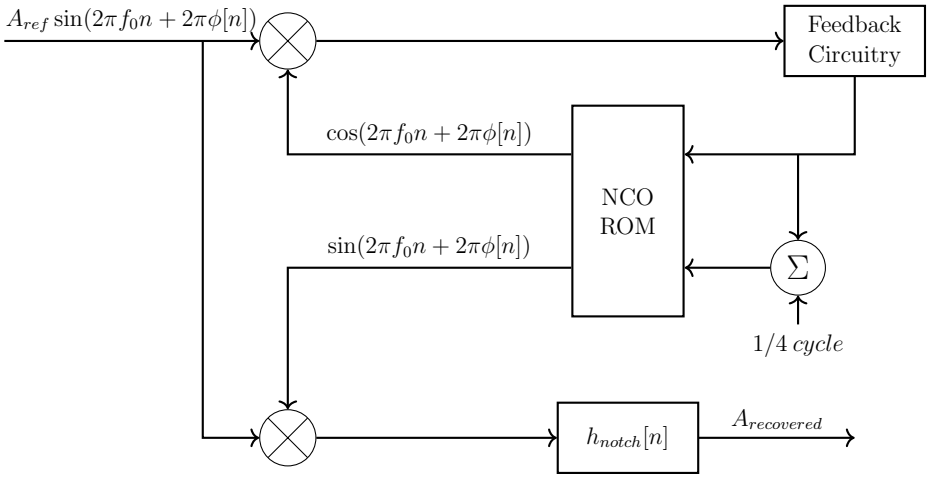

Figure 6.12 Amplitude Recovery Circuit

sufficiently low for the RAM on the FPGA to store 90 seconds of data. While the RAM is collecting the data the feedback loop inside the PLL is broken so that the phase changes in the network are not compensated by the control loop in the PLL. The circuit operation was described in detail in Chapter 4. The contents of the RAM are then exported to the PC via JTAG where they are analyzed using MATLAB.

The goal of the circuit was to be able to capture a large amount of phase fluctuation data and to be able to measure changes in phase as small as 300μ radians. As discussed in Chapter 5, due to the noise inherent to the ADC that is unachievable, but the circuit can come close. A plot showing the baseline noise at the output of the measurement circuit is shown in Figure 6.13

Figure 6.13 shows the noise floor of the phase measurement circuit to be at approximately 500μ radians, meaning that detection of a 300μ radian phase change is not achieved, but is near enough that the circuit can be considered as operational. The baseline noise shown here is for the DAC connected to the "out" coupled port on a hybrid coupler and the ADC connected to the "in". The combined port was connected to a modeled DOCSIS cable network shown in Chapter 7.

The transmit and receive sides of the circuit were connected digitally and the transmit

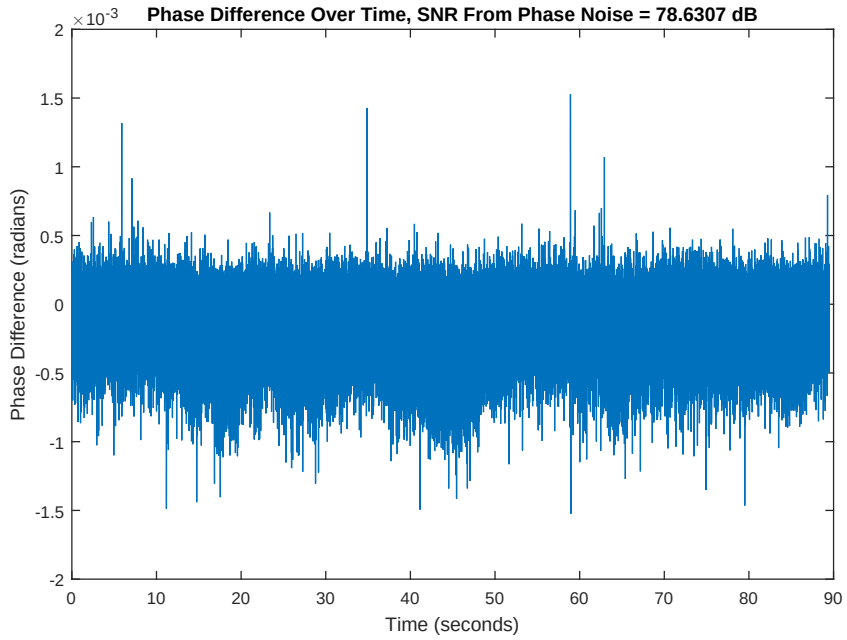

Figure 6.13 Baseline Noise at Output of Phase Measurement Circuit

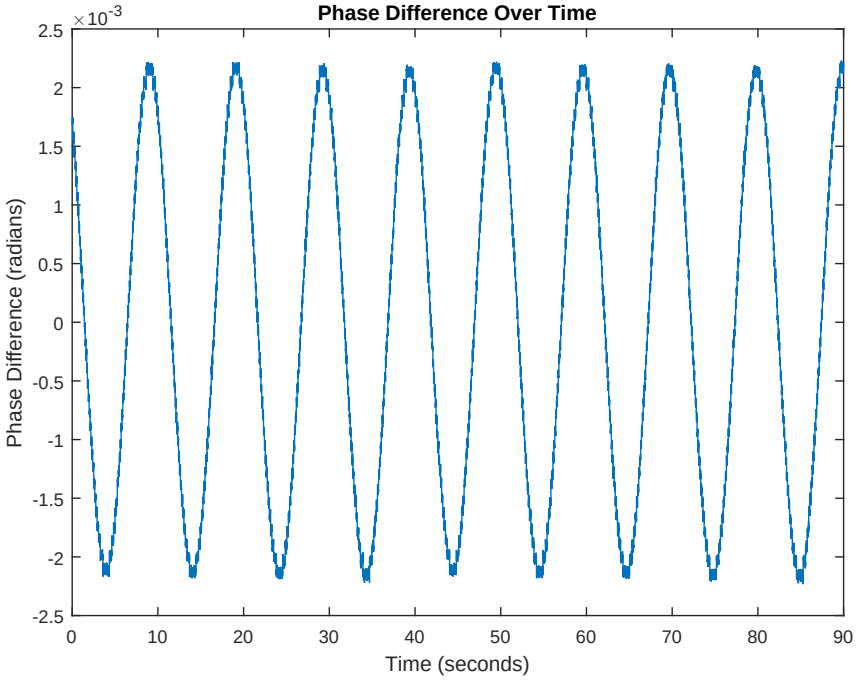

Figure 6.14 Hardware Phase Detector Output

sinusoid phase modulated with a 1/10 Hz sinewave of amplitude $\frac{35}{2^{17}} = 0.000267$ cycles. Due to the digital connection, the DAC/ADC was bypassed and a result the noise floor on the received signal is significantly lower than it would be otherwise. The sinusoid is shown to be properly detected at the filtered output of the phase detector in Figure 6.14.

7. Lab Measurements

In this chapter a series of experiments are performed on a model of a DOCSIS cable plant and the phase fluctuation data collected using the PLL measurement circuit. A model network is constructed to emulate a DOCSIS cable network to make the measurements of interest. A block diagram showing the different components in the model network is shown in Figure 7.1.

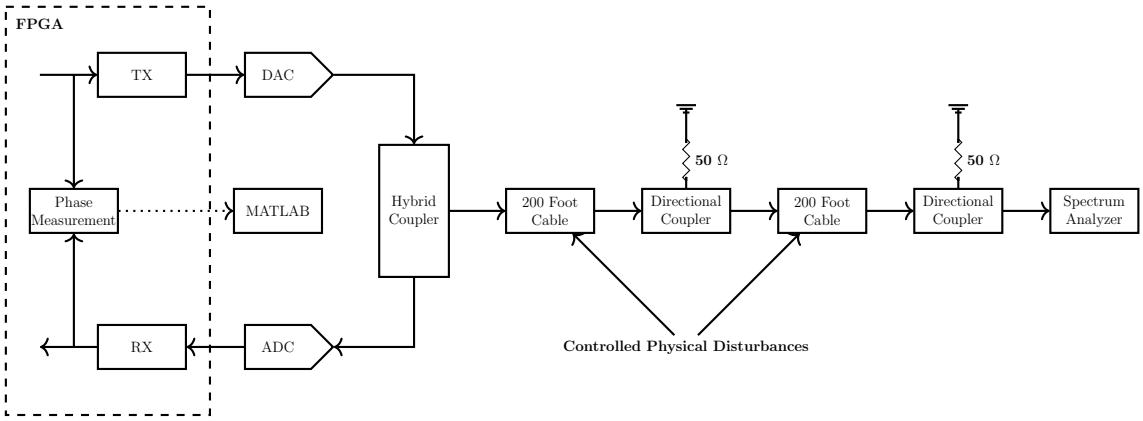

Figure 7.1 Lab Measurement Apparatus

The network model contains the key components of a DOCSIS network. The node is emulated by an FPGA which contains the transmitter, receiver, and phase measurement circuits needed for the experiments. Attached to the FPGA is a DAC/ADC board that performs the digital analog conversions needed. The DAC and ADC are then connected to the isolated ports of a 50Ω hybrid coupler. Attached to the combined port of the hybrid coupler is the rest of the network. The coupler is directly connected to a 200 foot 50Ω coaxial cable.

In a DOCSIS cable plant a characteristic impedance of 75Ω is standard and the input/output impedance of the ADC/DAC is 75Ω to reflect this. The cables and RF devices in the model network constructed for these experiments are all 50Ω. The mismatch in characteristic impedance between the DAC/ADC and the rest of the network means that a reflection coefficient of $\Gamma = \frac{50-75}{50+75} = -0.2$ is present at the 4-wire side of the hybrid coupler. This mismatch is unavoidable as the only components available were 50Ω devices.

Attached to the far end of the first 200 foot cable is a 20 dB directional coupler. The directional coupler is put in place of the taps used in a real DOCSIS network, which are just specialized directional couplers. The coupled port of the directional coupler is terminated with a 50Ω load so that it will not generate an echo. The out port of the directional coupler is then connected to another 200 foot cable which is connected to another directional coupler terminated in the same way. The out port of this directional coupler is then connected to a spectrum analyzer which is internally terminated so that it will not generate an echo.

The network model shown in Figure 7.1 shows two sections of 200 foot cable. There are two sections of cable present so that the model would better represent a small DOCSIS network, which would have many sections of cables and many taps. The cables used in these experiments all attenuate higher frequencies as discussed in Chapter 2. Due to the length of the cables used in the experiments the low pass filter effect is fairly significant. The combined loss through the entire network from the output of the hybrid to the spectrum analyzer at the test frequency of 189.75 MHz is approximately 42 dB. Due to this any echoes that are generated in the experiment at the far end of the network are greatly attenuated and will be far more difficult to measure. All experiments performed on the cable in this chapter are performed on the 200 foot length of cable nearer to the hybrid coupler to help mitigate this effect. If the experiments were performed on the far end length of cable the echoes would be much smaller in amplitude and any changes in phase would be smaller and harder to detect.

The base line noise of the phase measurement circuit connected to this network model is shown in Figure 6.13. The metric shown in the figure, SNR from phase noise, is a rough

idea of how much the phase fluctuations in the network lower the SNR of the received signal based on the equations derived in Chapter 3. The base line noise of the circuit is shown to be approximately 80 dB and any SNR in the received signal below this is considered to be caused by the fluctuations in phase caused by the experiment carried out to the network.

The SNR metric on its own is a poor describer of the phase fluctuations of the network as single large spikes in phase fluctuation will not significantly degrade the SNR. It is for this reason that the SNR along with the peak phase variation is used to evaluate the severity of the phase response variations. The smallest phase fluctuation that the system is able to measure is approximately 500 μradians. In the subsequent sections various experiments are carried out on the model of the network and their effect evaluated using these two metrics.

7.1 Effects of Physical Disturbances to Coaxial Cables

In the network used to connect the DOCSIS node to subscriber homes the coaxial cables used are often exposed to the elements. This can come from them being hung from utility poles or buried shallowly enough that they are compressed from people walking over them or cars driving above them. A typical cable line is shown hung from poles and going to a subscriber house hold in Figure 7.2.

The issue has been raised that something as simple as a strong wind may cause the cables to sway and shake enough to cause phase drift in the echoes that the echo canceler will not be able to track. To mimic this issue the cables in the model DOCSIS network shown in Figure 7.1 were put through a series of experiments to determine the effect on the phase response of the network. The resulting phase fluctuations from these experiments are shown in Figure 7.6, Figure 7.4, and Figure 7.8.

The first experiment on the effect of physical disturbances to the cable was to rhythmically flex the cable in an attempt to mimic the effect of a strong wind. The cables in the network were coiled into an ellipse with a major axis of 65 cm and a minor axis of 25 cm. The cable was coiled 40 times with each coil having a circumference of approximately 5 feet. The

Figure 7.2 Coaxial Cables Hung From Utility Poles

major and minor axis of the ellipse was rhythmically reversed to flex the cable, as shown in Figure 7.3. This experiment was carried out over a duration of 90 seconds and the effects are shown in Figure 7.4.

The phase difference shown in Figure 7.4 shows that there is a significant effect on the phase response of the network when the coaxial cables are rhythmically flexed. The SNR has degraded approximately 26 dB from the experiment and there is a peak phase fluctuation of approximately 9000 μradians. Each peak in the results shown correspond to the time when the positions of the cables were reversed, showing that a large shift in the response of the network can happen very quickly when the network is physically disturbed.

The next experiment carried out on the cables was to emulate the effect of the wind lightly swaying them in the breeze. The same coil of cable from the first experiment was hung up and lightly shaken over a period of 90 seconds. The hung cables are shown in Figure 7.5.

The effects of this experiment are shown in Figure 7.6. It can be seen that there are

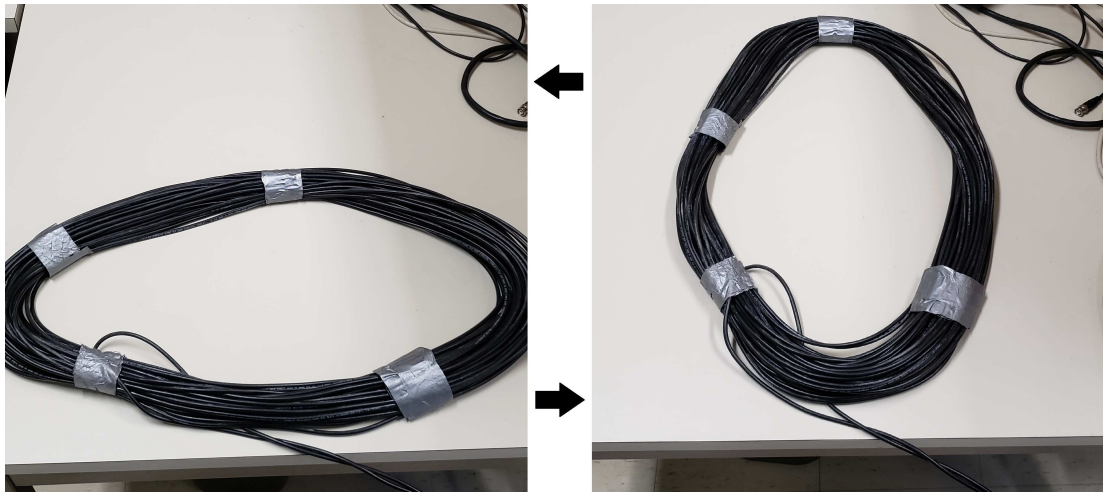

Figure 7.3 Experiment on Flexing Cable in a Model Network

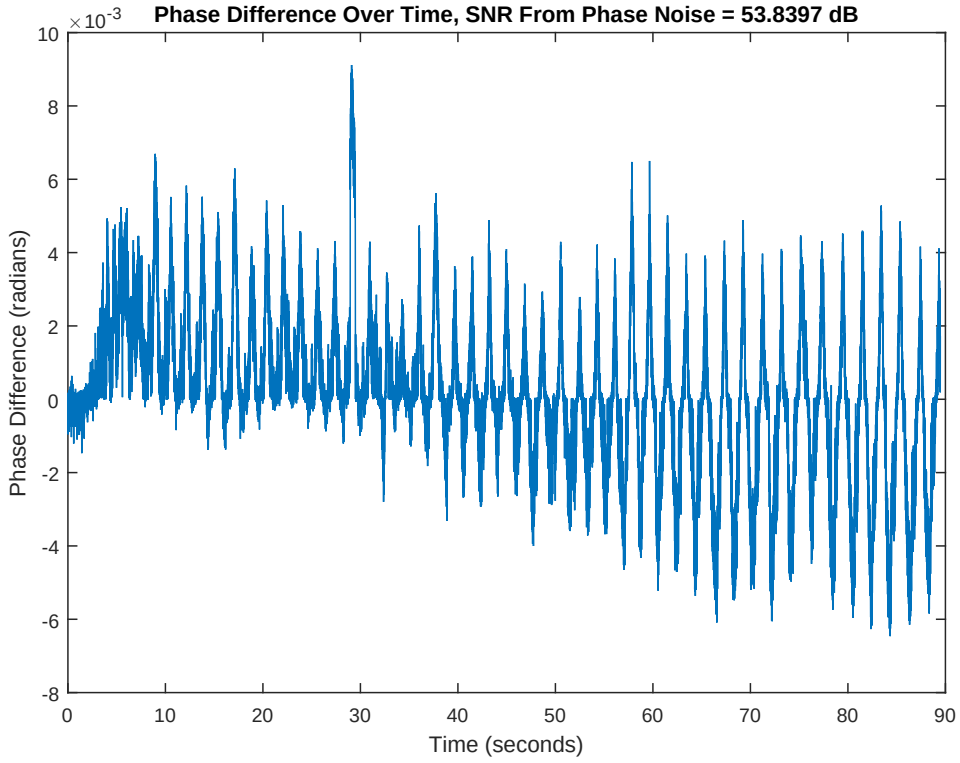

Figure 7.4 Phase Fluctuations from Flexed Cables

no strong peaks in the phase fluctuation as seen in Figure 7.4, as expected as the physical changes to the network were much less severe in this experiment. That being said there was

Figure 7.5 Hung Cable for Experiment on Lightly Shaking

still a SNR of degradation of 23.4 dB and a peak phase fluctuation of 8000 μradians. Due to the changes being much less sudden in this experiment, they are likely of less significance even though they appear comparable to the last experiment.

The final experiment carried out on the coaxial cables was to mimic the effect of cars and people moving over buried cables. The coil described in the previous two experiments was reshaped into a circle with a diameter of approximately 50 cm. The cable was then placed on the ground and a large piece of plywood placed on top of it to attempt to distribute the effect of the pressure more evenly over the cable. The plywood was then walked over in a circular pattern for 90 seconds and the data collected. A photo showing the plywood used in the experiment is shown in Figure 7.7.

The effects of this experiment are shown in Figure 7.8. Similar to the first experiment there are large sudden shifts in the phase response seen in this experiment. These shifts correspond to each step taken over the cables. The SNR degradation was approximately 26.5 dB and the peak phase fluctuation was approximately 8000 μradians. The sudden shifts in the phase response mean that this experiment would likely have a profound effect

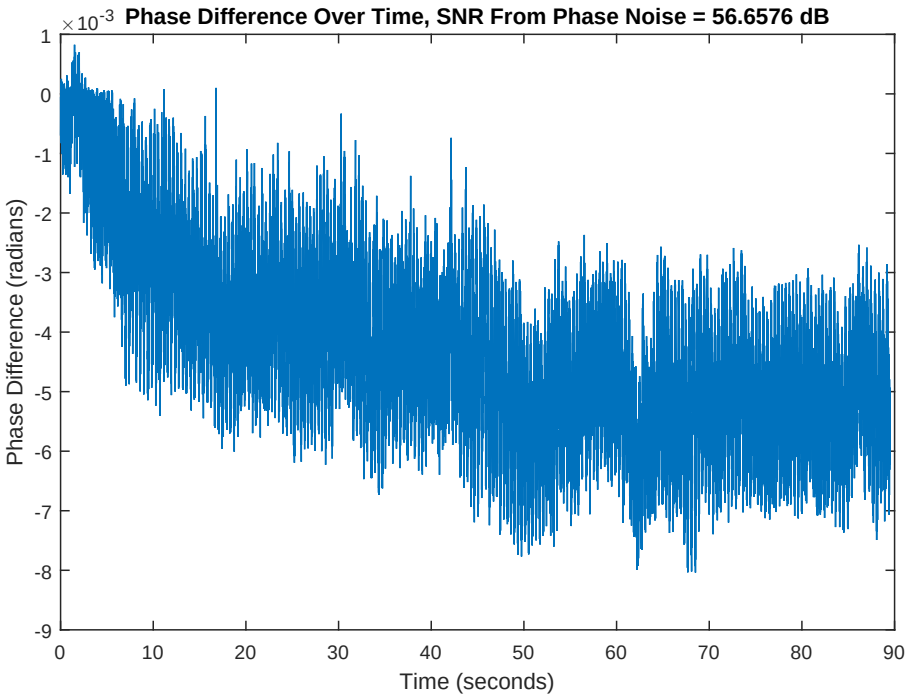

Figure 7.6 Phase Fluctuations from Shaken Cables

on the echo canceller as it would certainly struggle to track the response of the network.

7.2 Effect of Heated Cables

Concern has been raised that the effect of the sun rapidly heating cable plants exposed to the elements could rapidly shift the echo profile of the network. In desert environments such as the one in Phoenix Arizona the nights can become quite cold and the effect of the rising sun on hung cables such as the ones in Figure 7.2 can be significant. This rapid heating may lengthen the cable by increasing the sag between the poles and/or change the propagation characteristics of the cable thereby changing the characteristics of the network. It is possible that the rate of change is faster than the echo canceller is able to track.

The cable plant in the model network was again coiled and hung like in Figure 7.5 to attempt to mimic this phenomenon. A 1200W hair dryer was then placed 1 inch away from the hung cable coils and turned on for 90 seconds. During this time the phase variation

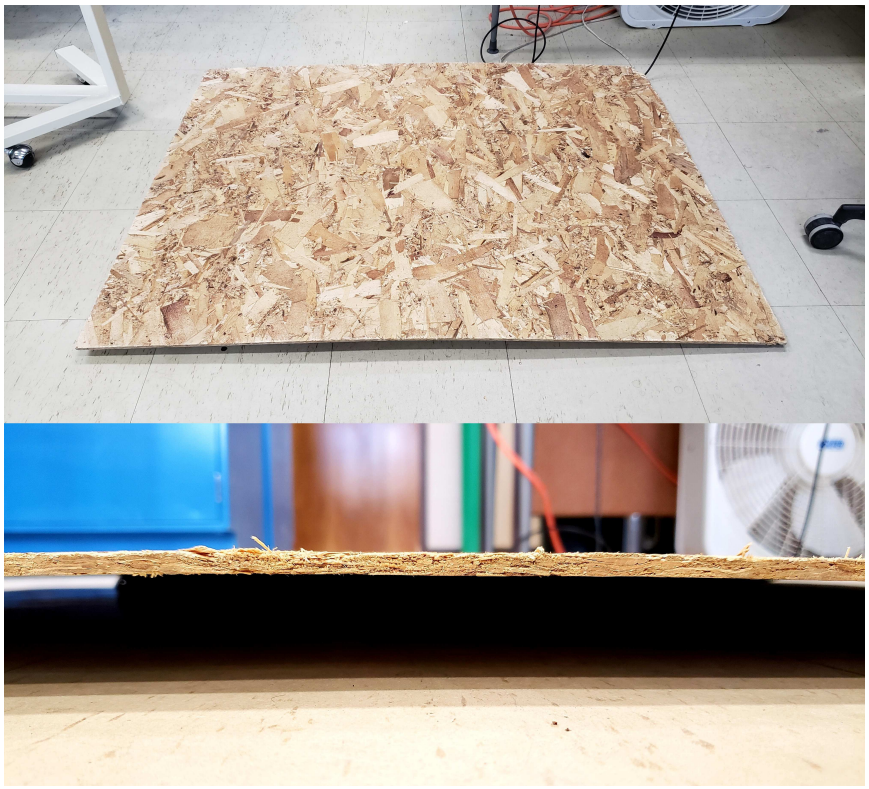

Figure 7.7 Walked Over Cable Experiment

data was collected. This experiment only heated one section of the cable plant but as the cable was coiled many times the effect was distributed throughout the length of the cable. If the entire length of the cable could be uniformly heated it would better represent the desired phenomenon. It is noted that any increased sag that may come from heating cables is not reflected in this experiment. A picture showing the experimental set up is shown in Figure 7.9.

The results of this experiment are shown in Figure 7.10. Approximately 40 seconds pass before any significant variation in the phase is seen. This is likely caused by the time it takes for the heat to transfer from the sheath of the cable to the core. Once the core begins to heat up the change in phase is rapid. The SNR is degraded by approximately 36 dB and there is a peak phase variation of 20000 μradians. The change in phase due the heating appears to be approximately linear once the core begins to heat with a slope of ≈ -400 μradians/second. It

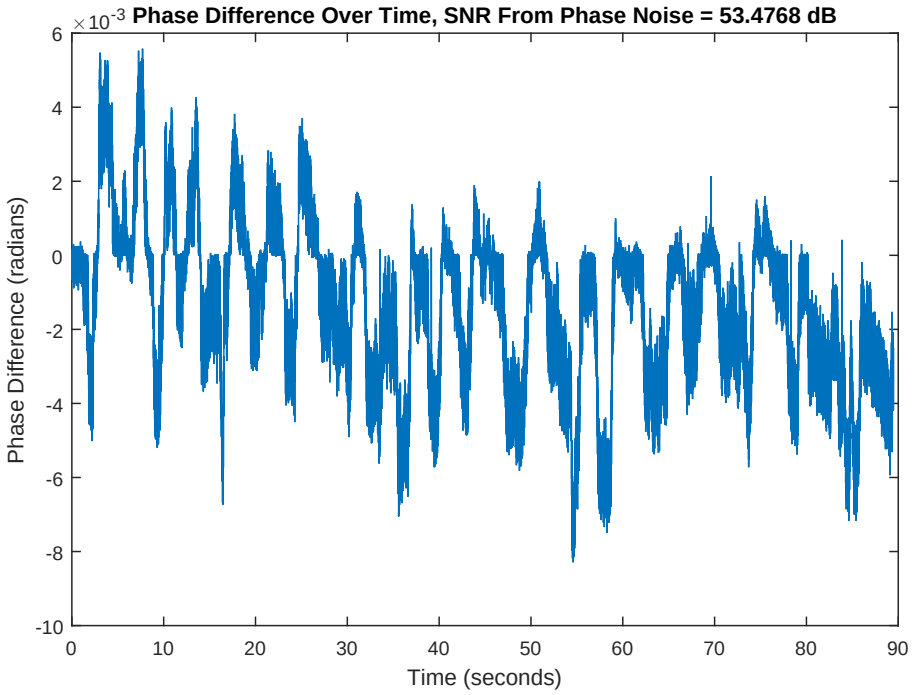

Figure 7.8 Phase Fluctuations from Walked Over Cables

is likely that this trend would slow down and stop once the core becomes as hot as the source heating it allows for. Given the apparent linearity of this phenomenon it is possible that the echo canceller may be able to track the changes in phase or perhaps be preprogrammed to predict for it.

7.3 Effect of Sudden Load Changes

In this experiment the model DOCSIS cable plant described in Figure 7.1 was subjected to sudden load changes and the resulting phase fluctuations recorded. The 50Ω terminations on the directional couplers were both quickly removed and replaced at regular intervals and left permanently open to simulate the sudden load changes. This has the effect of changing the reflection coefficient Γ_L described in Chapter 2. The taps used had a coupling loss of 20 dB. This means that when the tap's coupled port was left unterminated a echo 40 dB below the propagating signal was generated. This is because an unterminated load has a reflection

Figure 7.9 Heated Cable Experiment

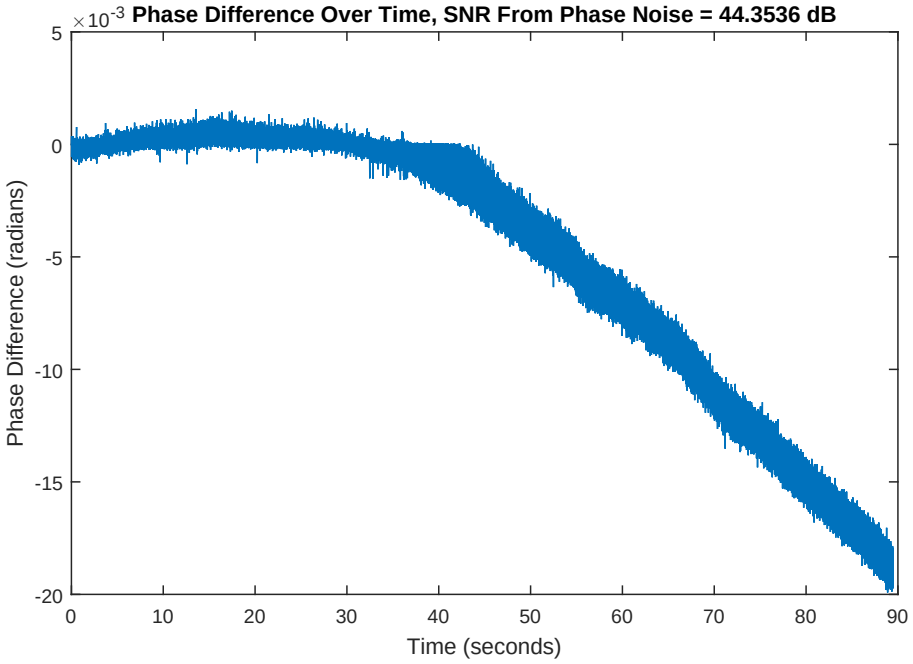

Figure 7.10 Phase Fluctuations from Heated Cables

coefficient of $\Gamma = 1$ that would go through the coupled port twice. A figure showing the tap with and without the termination is shown in Figure 7.11. With the changing reflection coefficient the echo profile of the entire network changes and the phase response subsequently rapidly shifts.

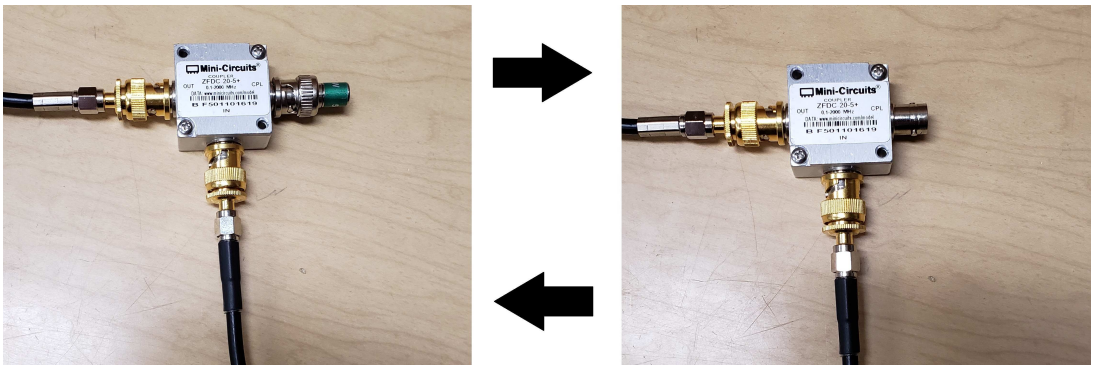

Figure 7.11 Terminated and Unterminated Tap Experiment

The phase locked loop measurement circuit records these changes in the phase response and exports them to MATLAB. This experiment was performed to determine the effect of users or technicians quickly shifting the loads connected to the taps in the network. This can come from quickly unplugging a cable modem, leaving a wire unconnected, or forgetting to properly terminate test points. The resulting phase fluctuations are shown in Figure 7.12 for a tap unterminated and in Figure 7.13 for a tap alternately terminated and unterminated in 10 second intervals.

The effect of the unterminated tap is that of a constant shift in the phase response of the network of approximately 3000 μradians. This is combined with a slightly higher noise floor likely caused by the introduction of more echoes and some weak pick up from the unterminated load. This sudden shift is immediately significant but as it mostly constant an echo canceller would quickly be able to adapt to the new network response and alleviate the issue. The alternately unterminated and terminated tap is much more problematic as if for some reason a load was rapidly switched on and off, such as in a faulty connection, the echo canceller would not be able to keep up and the noise floor would rise significantly.

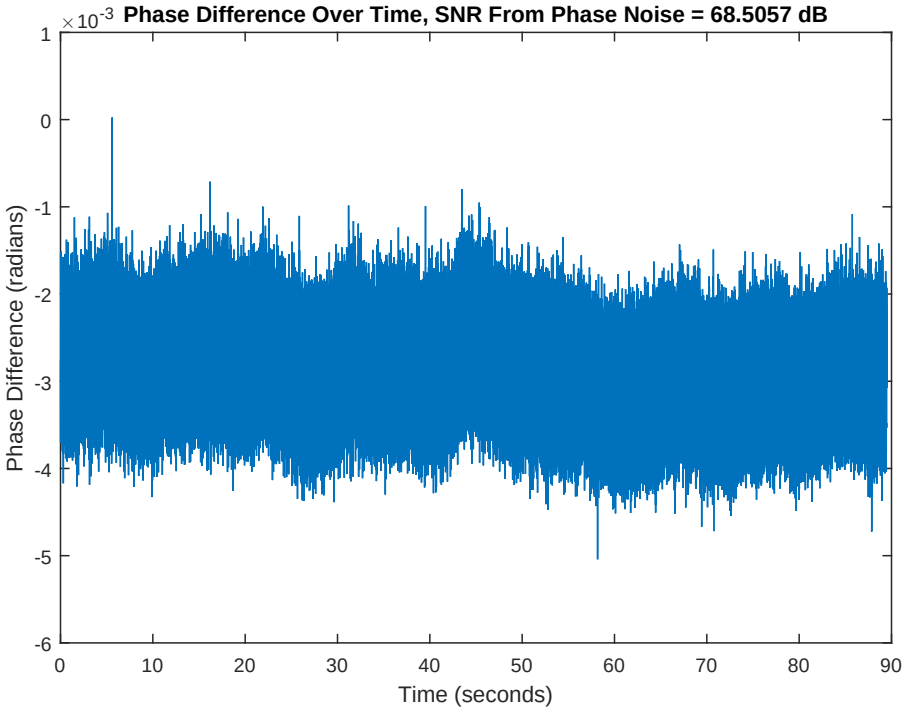

Figure 7.12 Phase Fluctuations from an Unterminated 20 dB Tap

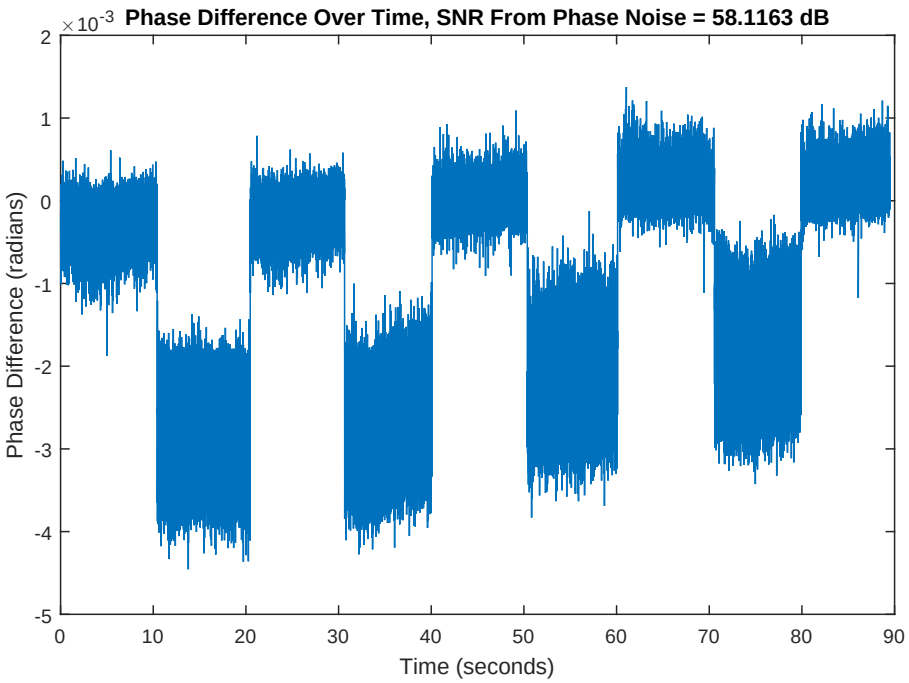

Figure 7.13 Phase Fluctuations from a Alternately Terminated and Unterminated 20 dB Tap

8. Summary and Future Work

8.1 Contributions and Results

The main contribution of this thesis is the development and implementation of a high performance circuit that can be used to measure the fluctuations in the phase response of a DOCSIS CATV network. This measurement circuit is implemented on a high speed FPGA and connects to the input and output of a cable network. This device measures the stability of the phase response of CATV networks with sufficient accuracy to determine if the network will support full duplex operation.

Data is collected in 90 second chunks in a .hex format that is easily converted to readable data by a MATLAB script. A JTAG interface connects the measurement circuit to a PC to allow data to be continuously exported. The JTAG interface also allows for the ROMs used to generate the probing signal to be loaded/updated with any arbitrary signal while the circuit is operating.

The measurement circuit is built around a phase locked loop that is designed to measure the fluctuations in phase of a sinusoid of a range of frequencies. The PLL works by locking the phase of a locally generated NCO to the phase of a reference sinusoid. Asserting the measurement "start" signal breaks the feedback of the loop while the difference between the generated and received sinusoid's phase is collected over 90 seconds. As the phase of the local sinusoid was locked to the phase of the reference at the time of the measurement start any deviations in the phase of the reference over the measurement time period are recorded. The JTAG interface and built in LUTs inside the PLL allow for the PLL to be tuned to a

wide range of frequencies.

A series of experiments were performed and their results recorded on a small scale DOCSIS network constructed in the lab. It was found that a number of physical disturbances to the network had the potential to rapidly shift the phase response of the channel. The effects of heating the cables was found to change the phase response of the channel rapidly and linearly with a peak phase shift of 20000 μradians. It was also determined that physical forces to the cables in the network had the propensity to rapidly change the phase response of the channel. When the cables were rhythmically flexed it was found that the phase response rapidly changed correspondingly. These shifts in phase response would certainly prove difficult for an echo canceller to track. The effects of improperly terminating the taps in the network were also explored. It was found that unterminated taps caused constant phase shifts in the network. These phase shifts would not affect the echo canceller unless they very suddenly changed, in which case the echo canceller performance would degrade until it adapted to the new phase offset.

In order to deal with these changes in phase an echo canceller will need to be able to adapt to a changing network. It was assumed in this thesis that the adaption time of an echo canceller would be at most 10 seconds. The results found in this thesis show that a 10 second adaption time may not be sufficient for some events. Sudden events such as when the coaxial cables are shaken and flexed change the phase faster than the canceller can track. More gradual effects such as heating the cable plant may be trackable by an echo canceller with a 10 second adaption time, but there would almost certainly be residual echoes. If the response of the network was in a constant state of flux it is possible than an echo canceller would never properly remove the interference on the received signal resulting in an unacceptable degradation of the noise floor.

8.2 Future Work

Due to a number of time constraints several concessions were made in the design and testing of the phase measurement circuit. It would be interesting to connect the phase measurement device to a real DOCSIS network in the field to determine the actual phase fluctuations in the network. Unfortunately, such a field test was not able to be set up in time for this thesis.

Finally, it would be interesting to have the phase measurement circuit run in tandem with an actual echo canceller to see how the phase fluctuations recorded by the phase measurement circuit actually manifested in uncancelled echoes. The theory for this was developed but to see the real world data would be interesting none the less.

Appendix Full Duplex Remaining System Components

The purpose of this appendix is to expand on some of the topics skipped over in Chapter 2. It should be read by readers looking for more background information on components of a DOCSIS network.

A.1 MAC Layer

The MAC (Media Access Control) layer of DOCSIS is responsible for controlling the physical layer. It deals with a number of key features of DOCSIS such as:

- Bandwidth allocation
- Providing upstream minislots
- Variable length packets
- Quality of Service (QoS), providing bandwidth and latency guarantees and creation, management and deletion of dynamic flows
- Range of data rates

The MAC layer of the DOCSIS network is essentially responsible for controlling all of the parameters of modulation and organizing the data that is being transmitted. The MAC layer is a vital part of the structure of DOCSIS but its functions do not immediately affect the topics researched in this thesis.

A.2 DOCSIS PHY Layer Components

A.2.1 Convergence Layer

The physical layer starts with a "convergence layer" that essentially acts as a receiver between the MAC layer and the PHY layer. The MAC layer generates packets of data from some source and the convergence layer receives them in the PHY layer.

A.2.2 Forward Error Correction

The Forward Error Correction (FEC) layer is another important part of the PHY layer. The FEC encoder adds some forward error correction to the packets received in the convergence layer. DOCSIS 3.1 FEC encoding uses a concatenated BCH-LDPC (Bose-Chaudhuri-Hocquenghem Low Density Parity Check) encoder followed by shuffling the bits in a code word around via bit interleaving [7]. The complexities of the FEC encoding are out of the scope of this research and as such will not be further discussed.

A.2.3 Constellation Mapper

The constellation mapper maps parallel streams of data into a sequence of quadrature constellations. A common form of constellation mapping is Quadrature Amplitude Modulation which divides the data into two streams and maps the bits in one stream to in-phase amplitudes and the bits in the other stream into quadrature amplitudes. After mapping the two streams they are treated together as one QAM symbol. An example of a 16-QAM constellation is shown in Figure A.1.

Figure A.1 shows a one-to-one mapping of a 4-bit binary word to an in-phase and quadrature amplitude for a 16-QAM symbol. For example, if the constellation mapper was to receive the bits (0001) it would map the output amplitude to $3 + j1$, which is an inphase amplitude of 3 and a quadrature amplitude of 1. The OFDMA modulator shown in Figure 2.2 has each of its subcarriers QAM modulated, it is not a requirement that each subcarrier have the same constellation mapping. In a DOCSIS 3.1 QAM constellation the subcarrier mod-

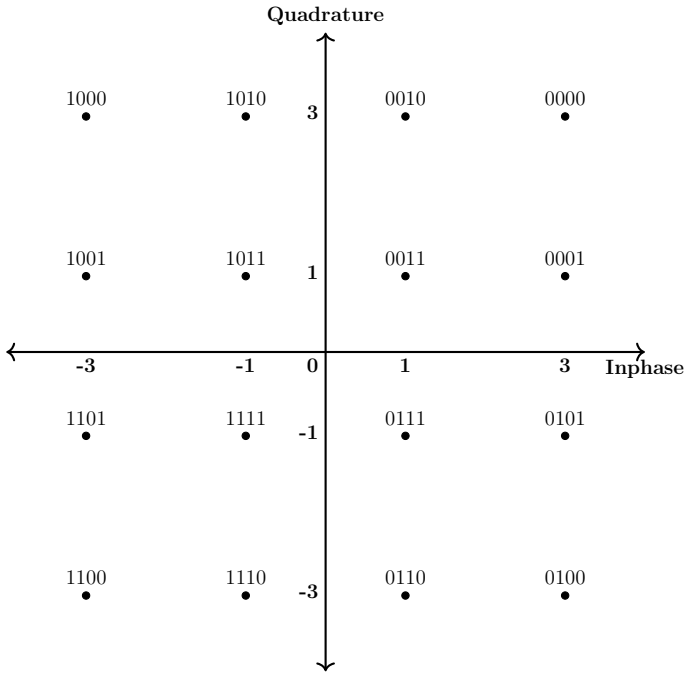

Figure A.1 Example Gray Coded 16-QAM Constellation for 4 Example Bits (y_1, y_2, y_3, y_4)

ulation schemes include zero bit loaded sub carriers, 16-QAM, 64-QAM, and all the way to 4096-QAM with both square shaped constellations like the one in Figure A.1 and non-square constellations. The CMs and CMTS also have the option to use 8192-QAM and 16384-QAM modulation orders [7].

The size of the QAM constellation directly effects the data rate of the system, with higher QAM orders allowing for higher data transmission rates. However, with higher QAM orders the constellation becomes more and more densely packed. This leads to problems in the receiver end as high order constellations are more prone to noise moving a signal from one constellation point to another leading to bit errors. For this reason the QAM constellation size and therefore the throughput of the system is limited by the noise in the system.

For QAM systems a metric known as the MER (Modulation Error Ratio) is a measure of the quality of the signal received from a QAM transmitter. The channel will add some noise to the transmitted symbol, this causes the received constellation point to move away from

the intended points plotted on the signal space shown in Figure A.1. As discussed in the above paragraph if this noise causes the received constellation point to move too far away from the original transmitted point it will cause a bit error. The MER is a measure of how close the received signal is to the originally transmitted signal.

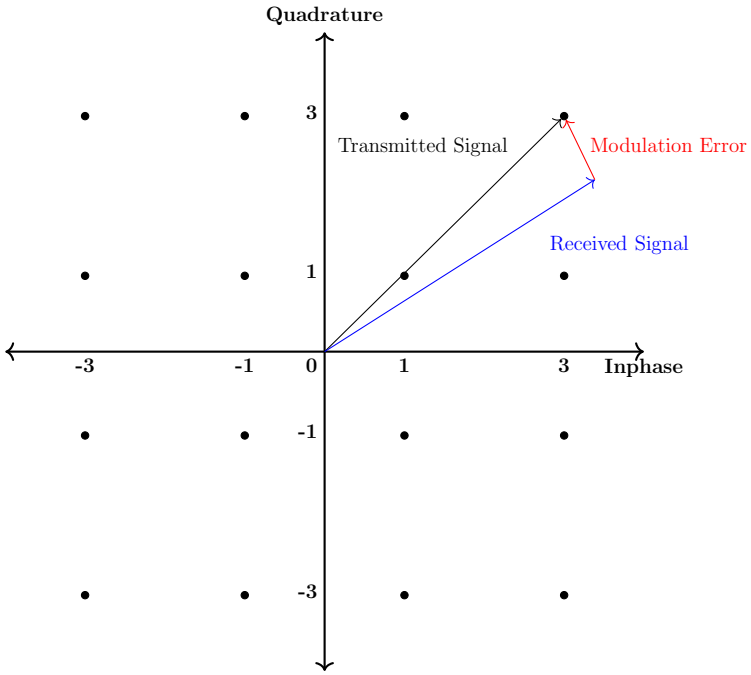

Figure A.2 Modulation Error

Mathematically the modulation error is the vector subtraction of the transmitted signal from the received signal. A more commonly used metric is the MER measured in dB with respect to the power in the target constellation point. The formula for MER is given in (A.1).

$$MER = 10\log_{10}\left(\left|\frac{\sum_{j=1}^{N}(I_j^2 + Q_j^2)}{\sum_{j=1}^{N}(\Delta I_j^2 + \Delta Q_j^2)}\right|\right) \tag{A.1}$$

where:

I = Real Part of Transmitted Symbol

Q = Quadrature Part of Transmitted Symbol

ΔI = Real Part of Difference Between Transmitted and Received Symbol

ΔQ = Imaginary Part of Difference Between Transmitted and Received Symbol

N = Number of Samples Observed

The formula for MER assumes that all constellation points are equally likely to be observed. To make a practical measurement, N must be large enough for all points to have occurred approximately the same number of times.

The minimum acceptable MER depends on the order of the constellation. As the QAM order gets larger the minimum acceptable MER also gets larger as there is less room in the signal space for error.

A.2.4 Inverse Fourier Transform

The IFFT block generates a discrete time domain signal. The operation of the IFFT is such that its input is a set of N complex amplitudes and its output is a set of orthogonal exponentials of the form $A_i e^{j2\pi in/N}$ for $i = 0 : N-1$. The output of the IFFT block is the sum of the N orthogonal exponentials organized into a length N group of samples. This length N group of samples is referred to as an OFDM symbol.

In the end, the IFFT block converts the frequency domain OFDM symbols into discrete time domain symbols suitable for transmission across the channel. The standard IFFT is described using Equation A.2 [35].

$$x[n] = \mathcal{F}^{-1}\{X[k]\} = \frac{1}{N_{FFT}} \sum_{k=0}^{N_{FFT}-1} X[k] e^{j\frac{2\pi nk}{N_{FFT}}},\ 0 \leq n \leq N_{FFT} - 1. \tag{A.2}$$

Where N_{FFT} is the size of the IFFT. DOCSIS 3.1 allows N_{FFT} to be either 4096 or 8092. The formula for the IFFT in DOCSIS is slightly modified from Equation A.2,

$$x[n] = \mathcal{F}^{-1}\{X[k]\} = \frac{1}{\sqrt{N_{FFT}}} \sum_{k=0}^{N_{FFT}-1} X[k] e^{j\frac{2\pi n(k-N_{FFT}/2)}{N_{FFT}}},\ 0 \leq n \leq N_{FFT}-1. \quad \text{(A.3)}$$

The DOCSIS IFFT equation differs in the scaling factor of $\sqrt{N_{FFT}}$ and the phase shift of $\frac{-N_{FFT}}{2}\ \frac{cycles}{sample}$. From this point forward the acronym "IFFT" refers to the DOCSIS defined IFFT in (A.3). The DOCSIS version of the IFFT is simply obtained from a standard IFFT by multiplying by a scaling factor.

$$\text{IFFT}_{\text{DOCSIS}} = \sqrt{N_{FFT}} \cdot e^{-j\pi n} \cdot \text{IFFT}_{\text{regular}},\ 0 \leq n \leq N_{FFT}-1 \quad \text{(A.4)}$$

At the output of the IFFT the sampling rate of the discrete time domain OFDM symbol is defined in the DOCSIS standard as 204.8 MHz. This corresponds to 25 MHz subcarrier spacing for a 8K IFFT and 50 MHz subcarrier spacing for a 4k IFFT. Each OFDM channel has a maximum bandwidth of 192 MHz corresponding to 3831 subcarriers in 4K mode and 7681 subcarriers in 8K mode.

If all the subcarriers in the OFDM symbol are used the OFDM symbol occupies 204.8 MHz, which exceeds the allowed bandwidth of 192 MHz of a OFDM channel. However, only the center 3800 subcarriers are used in 4K mode making the active bandwidth 192 MHz. In 8K mode only the center 7600 subcarriers are active yielding the same 192 MHz active bandwidth.

The full duplex band of the DOCSIS 3.1 Full Duplex spectrum is designed to occupy the region from 108 MHz to 684 MHz. This band is further defined into 3 192 MHz sub bands with each band containing one 192 MHz full duplex downstream channel and two 96 MHz full duplex upstream channels. The full duplex band is configurable to take up between 96 MHz and 576 MHz. Each bandwidth is always subdivided into 3 sub bands containing 1-2 full duplex upstream channels and 1 full duplex downstream channel. The spectral allocation is controlled by the MAC layer [7]. An example spectral allocation for a 576 MHz full duplex band is shown in Figure A.3.

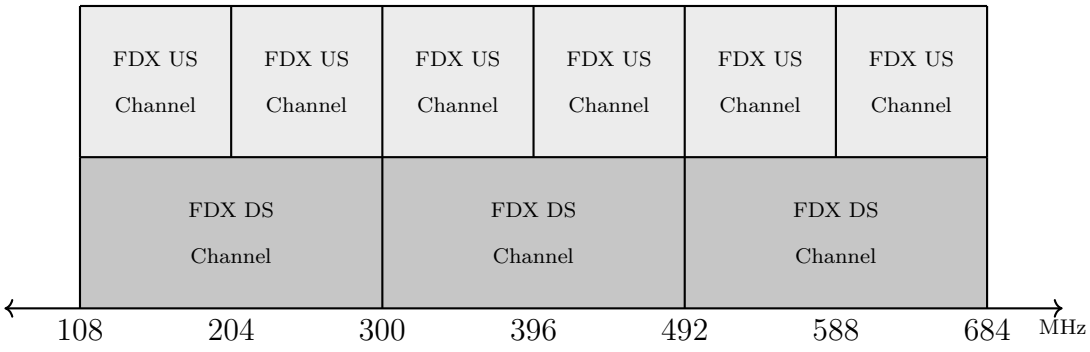

Figure A.3 Full Duplex Band Allocation

A.2.5 Cyclic Prefix and Roll Off Insertion

For each OFDM symbol a segment at the end of the symbol is appended to the beginning of the symbol. This is called the cyclic prefix and it is used to combat Inter Symbol Interference (ISI) caused by echoes in the channel. In DOCSIS 3.1 the length of the cyclic prefix, N_{CP}, has 5 possible values that are chosen based on the delay spread of the channel with a longer delay spread requiring a longer cyclic prefix. As adding the cyclic prefix lowers throughput it is desirable to keep it as short as possible.

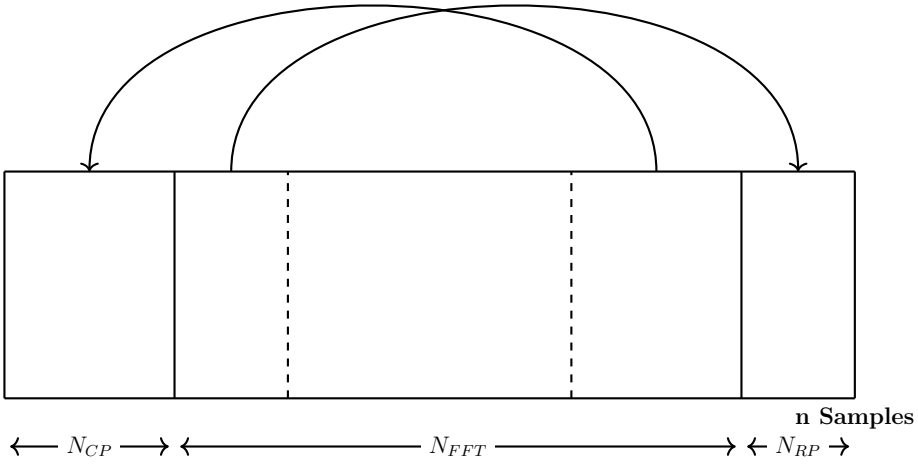

Figure A.4 OFDM Cyclic Prefix and Window Algorithm

For the purposes of windowing each symbol another segment at the beginning of the

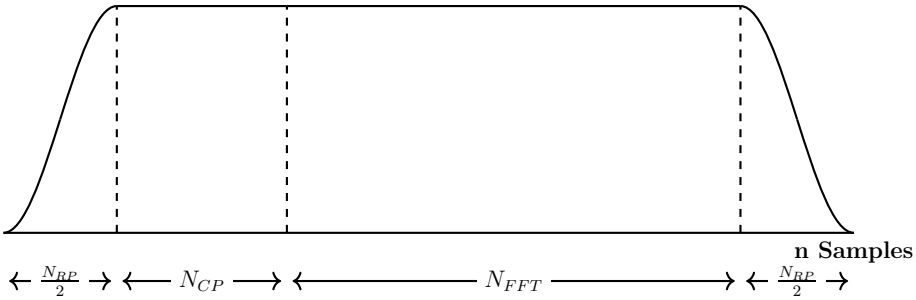

Figure A.5 Cyclic Prefix and Roll Off Insertion

symbol is appended to the end of the OFDM symbol. This segment is called a roll off period and it again has 5 possible lengths. It is referred to as N_{RP} in DOCSIS 3.1. The roll off period is windowed using Equation A.5. As the window is symmetric, only the right half of the window needs to be defined. The total length of the symbol at the output of the window module is defined as $N = N_{FFT} + N_{CP} + N_{RP}$.

$$w[\tfrac{N}{2}+n] = \begin{cases} 1, & for\ n = 0, 1, \ldots, (\frac{N-2N_{RP}}{2}-1) \\ \frac{1}{2}(1 - sin(\frac{\pi}{\alpha(N-N_{RP})}(n - \frac{N-N_{RP}}{2} + \frac{1}{2}))) & for\ n = (\frac{N-2N_{RP}}{2}), \ldots, (\frac{N}{2}-1) \end{cases} \tag{A.5}$$

A larger roll off period will produce a sharper spectrum of the OFDM symbol. The trade off being that a larger roll off period causes each symbol to have a longer duration and as such lowers throughput [7].

A.2.6 Upstream PHY Layer

The upstream layer inside the node essentially does the reverse of the downstream layer. Instead of encoding, mapping, and modulating packets of data to be transmitted it demodulates, demaps, and decodes packets of data that are received in the upstream direction. Instead of an IFFT to convert the frequency domain OFDM samples to discrete time domain a FFT is used to convert discrete time domain OFDM samples into frequency domain samples. Instead of a QAM mapper to convert bits into signal levels a QAM slicer is used to convert signal levels into bits. The bits are then decoded using FEC algorithms and sent to

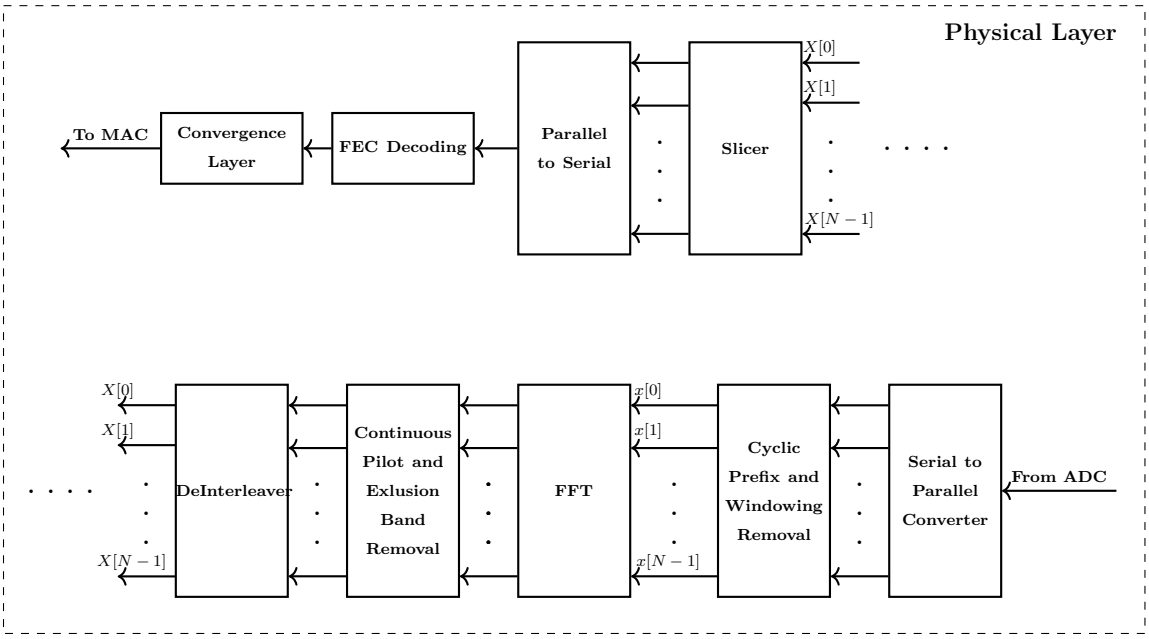

Figure A.6 Simplified Upstream PHY Processing

the MAC layer and to the CMTS. The overall process of the upstream PHY layer is shown in Figure A.6 [7]. As much of the operations in the receiver side of the PHY layer are simply mirrors of the transmitter side, steps not immediately relevant to the system level operation of the full duplex network will be ignored.

FFT

The receiver side of the full duplex node used an FFT to convert discrete time domain samples to frequency domain samples as apposed to the IFFT used on the transmit side. DOCSIS again used a modified version of the standard FFT.

$$X[k] = \mathcal{F}\{x[n]\} = \frac{1}{N_{FFT}} \sum_{n=0}^{N_{FFT}-1} x[n] e^{-j\frac{2\pi nk}{N_{FFT}}},\ 0 \leq n \leq N_{FFT} - 1 \tag{A.6}$$

The DOCSIS FFT again differs from the standard FFT in a phase shift and scaling factor.

$$X[k] = \mathcal{F}\{x[n]\} = \frac{1}{\sqrt{N_{FFT}}} \sum_{n=0}^{N_{FFT}-1} x[n] e^{-j\frac{2\pi n(k-N_{FFT}/2)}{N_{FFT}}},\ 0 \leq n \leq N_{FFT}-1 \qquad \text{(A.7)}$$

From this point on any references to the FFT will refer to the DOCSIS FFT. Using an IFFT and FFT for OFDM modulation is standard practice and leads to OFDM having high spectral and hardware efficiency.

QAM Slicer

After the received signal is passed through an FFT and is deinterleaved in both time and frequency it is passed into a QAM Slicer to convert the quadrature signal levels into sequences of received bits. A QAM slicer works by separating the signal space into so called decision regions and then examining which region the received signal falls into. For example if the received signal had inphase amplitude -2.3 and quadrature amplitude -3.7 a 16-QAM slicer would interpret that as "1100". Having decision regions instead of specific points allows the received signals to move away from their original transmitted amplitudes and still be correctly interpreted. As discussed in the section on QAM mappers the higher the QAM order the smaller that the decision regions become and the less room for error is acceptable in receiver. Modern DOCSIS systems have equalizers and pre-equalizers that correct for much of the channel response and greatly reduce the error on the received QAM-symbols. Even with the equalizers present the MER of the incoming signal will never be perfect. Reflections and micro reflections will degrade the MER of a signal as well noise sources such as burst noise, ingress noise, and phase noise. Phase changes in the signal caused by physical changes to the cable plant also have the capability to reduce the MER of the received signal greatly.

A.3 Taps

The taps shown inside the channel in Figure 2.1 are also directional couplers. The taps are designed to couple some signal from the "input port" to the "coupled ports" and pass the remaining power to the "through port". The couplers are designed in such a way that there

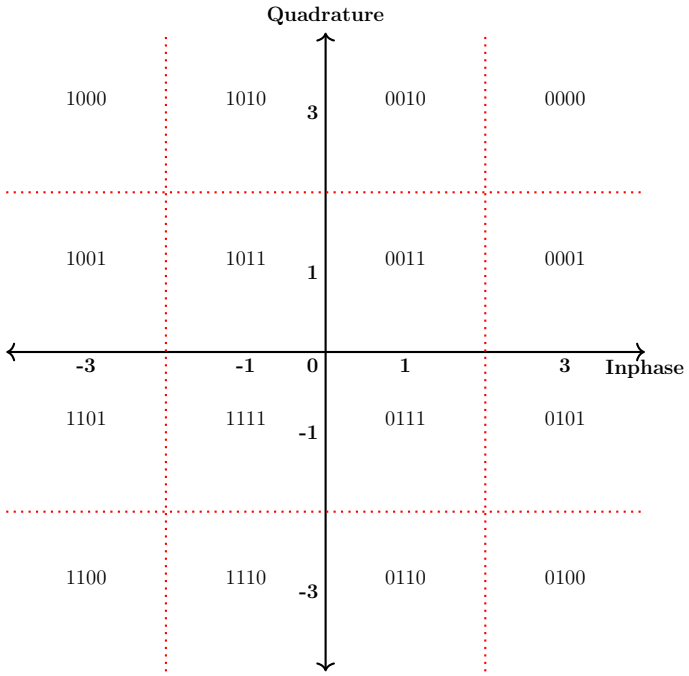

Figure A.7 16-QAM Slicer

is a great deal of isolation between the through port and the coupled port so that signals traveling in the upstream direction do not leak between the through and coupled ports. A tap is a directional coupler meaning it, like a hybrid, is truly a 4-port device but it is typically drawn as a 3-port device. This is because one of the ports is internally terminated. The directional property of directional couplers arises from the isolation between the coupled and through ports. Signals entering the through port, which means they are traveling in the upstream direction, are greatly isolated from the coupled port. Signals entering the input port, which mean they are traveling in the downstream direction, have a defined loss. In essence a percentage of the power is "tapped" off explaining the nomenclature of taps. The measure of a directional couplers directional properties is called directivity and is the difference between the isolation in dB and the coupled loss in dB.

$$\text{Directivity (dB)} = \text{Isolation (dB)} - \text{Coupled Loss (dB)} \tag{A.8}$$

A directivity of 30dB is considered good for most directional couplers [36]. A diagram

showing the power flow through a directional coupler is shown in Figure A.8. The figure shows the losses between each port.

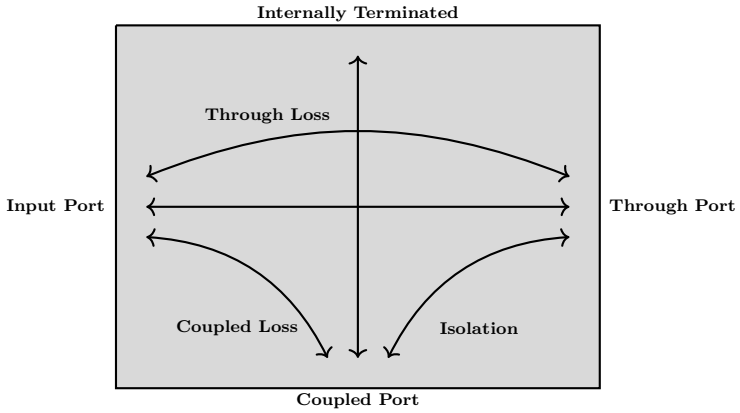

Figure A.8 Power Flow Through a Directional Coupler

In a CATV network the loss is chosen to provide the subscriber with a power level within almost 3dBm. The tap loss for subscribers down the chain is less as there is less power entering the coupled ports. Each tap in the CATV networks function is to provide some percentage of the power in the signal coming from the head end to some number of CMs. Most taps used in the CATV network have multiple coupled ports with 2-way, 4-way, and 8-way taps being commonly used in the CATV network. As the signal coming from the head end propagates through the channel it becomes attenuated and loses power to each tap. To counter act this each tap in the coaxial cable network will have progressively lower coupling losses so that the power going to the cable modems remains constant. For this reason each network is custom designed. An example network is shown in Figure A.9 highlighting the signal level as it propagates through the network.

The signal can be seen to be losing powers as it propagates through the network. The designer of the network carefully chooses the taps so that each CM receives approximately the same amount of power. At the end of the network there is a terminator of value Z_0, this is to prevent reflections back towards the node. For a CATV network the characteristic impedance is defined as 75Ω [7]. In practice it is not likely to have a perfect terminator and some echoes will be introduced. Each of the taps will also have imperfect terminations and

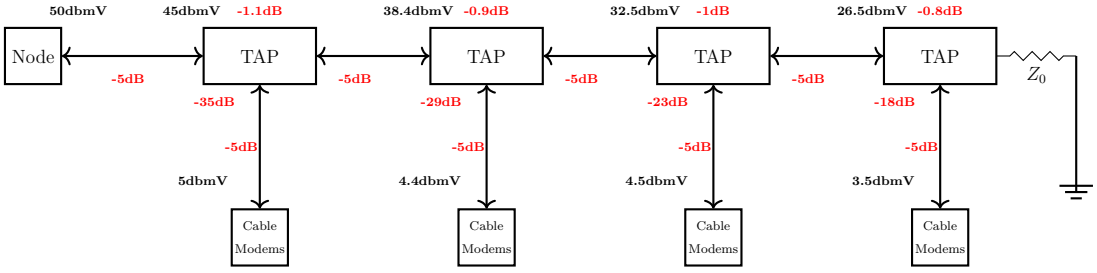

Figure A.9 Example Losses Through a Typical Coaxial Cable Network, Losses Are Shown in Red and Signal Levels in Black

introduce some echoes in the reverse direction. These echoes contribute to the reduction of the SNR on the signal propagating from the CMs to the node.

A.4 Cable Modems

The operation of cable modems inside the full duplex DOCSIS network is immensely complex. The cable modems job is to receive data from the network and then transmit data generated in the subscriber household to said network. At the most simple level the CMs act as small nodes inside the subscriber households. Each having its own transmit and receive path much like the node does.

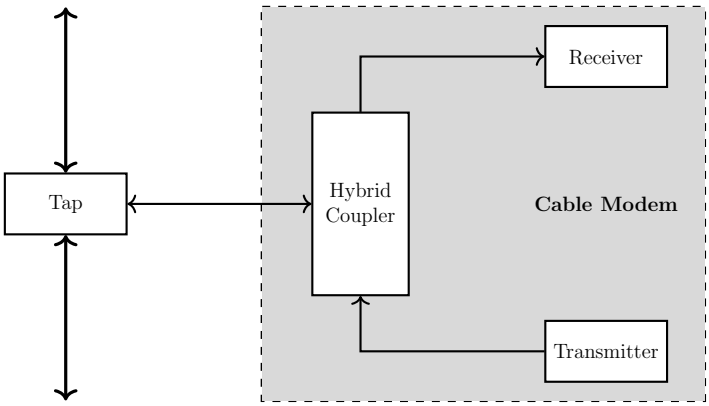

Figure A.10 Cable Modem Simplified Block Diagram

The cable network contains a large amount of these cable modems which all contribute to the upstream signal. There are often multiple cable modems attached to each tap meaning

that on top of the self interference from the transmitter to the receiver in the node the different cable modems will also start to interfere with each other. These groups of cable modems are called interference groups and the MAC controls them in such a way that the interference between them is minimized [7]. For the purposes of this thesis the intricacies of the cable modem transmission and receive algorithms will be ignored and the cable modems will be treated as a monolith that together create the upstream signal that the node receives.

References

[1] E. S. Smith, "The emergence of cable TV: A look at the evoloution of a revolution," in *IEEE, vol. 58*, pp. 967–982, July 1970.

[2] W. S. Ciciora, "An introduction to cable television in the united states," *IEEE LCS*, vol. 1, pp. 19–25, Feb 1990.

[3] J. S. Light, "Before the internet, there was cable," *IEEE Annals of the History of Computing*, vol. 25, pp. 96–95, April 2003.

[4] V. Brugliera, "History of compatibility between cable systems and receivers," in *IEEE International Conference on Consumer Electronics*, pp. 186–187, June 1994.

[5] Cable Labs, *Cable Labs Mission Statement*, 1988.

[6] Cable Labs, *DOCSIS 1.0 Radio Frequency Interface Specification, SP-RFI-C01-011119*, 1996.

[7] Cable Labs, *DOCSIS 3.1 Physical Layer Specification, CM-SP-PHYv3.1-I12-171026*, Oct. 2017.

[8] J. T. Chapman, "DOCSIS Remote Phy," White Paper, Society of Cable Telecommunications Engineers, Oct. 2013.

[9] W. Y. Chen and K. Kerpez, "Coaxial cable distribution plant performance simulation for interactive multimedia tv," in *Global Telecommunications Conference, 1995. GLOBECOM '95., IEEE*, vol. 1, pp. 173–177 vol.1, Nov 1995.

[10] A. Sahai, G. Patel, C. Dick, and A. Sabharwal, "Understanding the Impact of Phase Noise on Active Cancellation in Wireless Full-Duplex," in *Signals, Systems and Computers (ASILOMAR) Conference*, (Pacific Grove, CA, USA), pp. 29–33, Nov. 2012.

[11] L. C. II, *Digital and Analog Communication Systems*. Macmillan Publishing Company, 2 ed., 1987. ch. 2, pp. 34-35.

[12] D. M. Pozar, *Microwave Engineering*. Wiley, 4 ed., 2012. ch. 7, pp. 343-356.

[13] M. Cazarotto and R. d. R. Lopes, "Characterising micro-reflections in coaxial cable network," *IET Science, Measurement Technology*, vol. 11, no. 8, pp. 1043–1048, 2017.

[14] M. M. Driscoll and J. B. Donovan, "Vibration-induced phase noise: It isn't just about the oscillator," in *2007 IEEE International Frequency Control Symposium Joint with the 21st European Frequency and Time Forum*, pp. 535–540, May 2007.

[15] H. H. Nguyen and E. Shwedyk, *A First Course in Digital Communications*. Cambridge University Press, 2009.

[16] E. Eweda, "Stabilization of high-order stochastic gradient adaptive filtering algorithms," *IEEE Transactions on Signal Processing*, vol. 65, pp. 3948–3959, Aug 2017.

[17] S. Haykin, *Adaptive Filter Theory*. Prentice-Hall, Inc., 1991.

[18] Z. Andalibi, B. Berscheid, E. Salt, and H. H. Nguyen, "A fast-converging equalizer for upstream docsis channels," *IEEE Transactions on Broadcasting*, vol. 56, pp. 311–320, Sept 2010.

[19] M. Adams and V. K. Bhargava, "Use of the Recursive Least Squares Filter for Self Iinterference Channel Estimation," in *Vehicular Technology Conference (VTC-Fall), 2016 IEEE 84th*, (Montreal, QC, Canada), Sept. 2016.

[20] J. Stewart, *Calculus*. Brooks/Cole Publishing Company, 3 ed., 1995. ch. 10, pp. 655-666.

[21] A. Hati, C. W. Nelson, and D. A. Howe, "Effect of vibration on pm and am noise of oscillatory and non-oscillatory components at 10 ghz," in *Frequency Control Symposium, 2009 Joint with the 22nd European Frequency and Time forum. IEEE International*, pp. 524–529, IEEE, 2009.

[22] Agilent, *Network Analyzer Basics*, 2005.

[23] A. Pokatilov, T. Kübarsepp, and V. Vabson, "Effect of keysight 3458a jitter on precision of phase difference measurement," *IEEE Transactions on Instrumentation and Measurement*, vol. 65, pp. 2595–2600, Nov 2016.

[24] H. Yang, Y. Tu, H. Zhang, and K. Yang, "A Hilbert Transform based method for dynamic phase difference measurement," in *Control and Decision Conference (CCDC),2012 24th Chinese*, (Taiyuan, China), May 2012.

[25] J. Grove, J. Hein, J. Retta, P. Schweiger, W. Solbrig, and S. R. Stein, "Direct-digital phase-noise measurement," in *Proceedings of the 2004 IEEE International Frequency Control Symposium and Exposition, 2004.*, pp. 287–291, Aug 2004.

[26] F. M. Gardner, *Phaselock Techniques*. Wiley-Interscience, 3 ed., 2005. ch. 4, pp. 65-96.

[27] D. Richman, "Color-carrier reference phase synchronization accuracy in ntsc color television," *Proceedings of the IRE*, vol. 42, pp. 106–133, Jan 1954.

[28] G. A. Leonov, N. V. Kuznetsov, M. V. Yuldashev, and R. V. Yuldashev, "Hold-in, pull-in, and lock-in ranges of pll circuits: rigorous mathematical definitions and limitations of classical theory," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 62, no. 10, pp. 2454–2464, 2015.

[29] A. Papoulis, *Probability, Random Variables and Stochastic Processes*. McGraw Hill Education, 4 ed., 2002.

[30] W. Bennett, "Spectra of quantized signals," in *Bell System Technical Journal*, vol. 27, pp. 446–472, July 1948.

[31] G. Popek and M. Kampkic, "Low-spur numerically controlled oscillator using taylor series approximation," in *XI International PHD Workshop*, Silesian University of Technology, Poland, 2009.

[32] L. P. Sabel and F. J. Harris, "Using the high order nyquist zones in the design of efficient multi-channel digital upconverters," in *Personal, Indoor and Mobile Radio Communications, 1997. Waves of the Year 2000. PIMRC '97., The 8th IEEE International Symposium on*, vol. 2, pp. 355–360 vol.2, Sep 1997.

[33] fredric harris, *Multirate Signal Processing For Communication Systems*. Pearson Education Inc., 1 ed., 2004.

[34] I. H. H. Jorgensen, P. Pracný, and E. Bruun, "Hardware-efficient implementation of half-band iir filter for interpolation and decimation," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 60, pp. 892–896, Dec 2013.

[35] A. V. Oppenheim and R. W. Schafer, *Discrete Time Signal Processing*. Pearson International, 3 ed., 2010. ch. 9, pp. 716-762.

[36] T. Mukaihata, M. F. Bottjer, and H. J. Tondreau, "Rapid broad-band directional coupler directivity measurements," *IRE Transactions on Instrumentation*, vol. I-9, pp. 196–202, Sept 1960.